



(12) **United States Patent**  
**Oetting et al.**

(10) **Patent No.:** **US 10,491,705 B2**  
(45) **Date of Patent:** **Nov. 26, 2019**

(54) **VISUALIZATION FOR NETWORK  
VIRTUALIZATION PLATFORM**

(56) **References Cited**

(71) Applicant: **AT&T Intellectual Property I, L.P.**,  
Atlanta, GA (US)

(72) Inventors: **John Oetting**, Zionsville, PA (US);  
**Roger Wickes**, Roswell, GA (US)

(73) Assignee: **AT&T Intellectual Property I, L.P.**,  
Atlanta, GA (US)

(\* ) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 360 days.

U.S. PATENT DOCUMENTS

7,024,419 B1	4/2006	Klenk et al.	
7,224,362 B2	5/2007	Kincaid et al.	
7,454,439 B1	11/2008	Gansner et al.	
7,523,187 B1 *	4/2009	Lavallee	H04L 41/12 709/220
8,019,865 B2	9/2011	Chang et al.	
8,782,265 B1	7/2014	Bokotey	
8,935,396 B2	1/2015	Bokotey	
8,984,121 B1 *	3/2015	Labat	G06F 8/70 707/769
9,268,460 B1 *	2/2016	Xu	G06F 3/04815
9,418,455 B1 *	8/2016	Wong	G06T 11/206
9,647,909 B2 *	5/2017	Kuan	H04L 43/045

(Continued)

(21) Appl. No.: **14/847,573**

(22) Filed: **Sep. 8, 2015**

(65) **Prior Publication Data**

US 2017/0070594 A1 Mar. 9, 2017

(51) **Int. Cl.**  
**G06F 15/16** (2006.01)  
**H04L 29/08** (2006.01)  
**H04L 12/24** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **H04L 67/32** (2013.01); **H04L 41/12**  
(2013.01); **H04L 41/22** (2013.01); **H04L**  
**41/5058** (2013.01)

(58) **Field of Classification Search**  
CPC ..... H04L 41/12; H04L 67/1097; H04L 41/22;  
H04L 43/045; H04L 63/20; H04L  
41/0806; G06F 3/0484; G06F 3/04815;  
G06F 8/70; G06F 11/32; G06F 11/324;  
G06F 3/0481; G06F 9/505; G06T 11/206

See application file for complete search history.

OTHER PUBLICATIONS

Shneiderman, Ben, "Network Visualization by Semantic Sub-  
strates," IEEE Transactions on Visualization and Computer Graph-  
ics, Sep./Oct. 2006, pp. 733-740, vol. 12, No. 5.

(Continued)

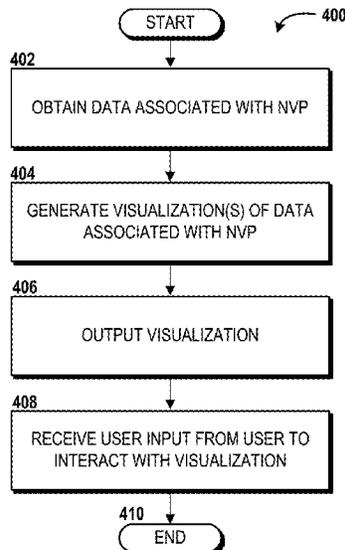
*Primary Examiner* — Dustin Nguyen

(74) *Attorney, Agent, or Firm* — Hartman & Citrin LLC

(57) **ABSTRACT**

Concepts and technologies disclosed herein are directed to  
visualizations for network virtualization platforms  
("NVPs"). According to one aspect disclosed herein, a  
computer system executing a visualization component can  
obtain data associated with an NVP. The computer system  
can generate a visualization of the data associated with the  
NVP. The visualization can include a three-dimensional  
central cube visualization that includes a plurality of smaller  
cube virtualizations, each of which includes a node repre-  
sentative of an application executable by at least a portion of  
a plurality of resources of the NVP. The computer system  
can output the visualization.

**20 Claims, 11 Drawing Sheets**



(56)

**References Cited**

U.S. PATENT DOCUMENTS

2001/0042118 A1\* 11/2001 Miyake ..... H04L 41/12  
709/223  
2005/0010386 A1\* 1/2005 Tharp ..... H04L 41/0806  
703/13  
2008/0077868 A1\* 3/2008 Bartucca ..... H04L 41/22  
715/736  
2008/0123586 A1\* 5/2008 Manser ..... H04L 41/12  
370/328  
2008/0281912 A1\* 11/2008 Dillenberger ..... G06F 11/324  
709/204  
2010/0110932 A1\* 5/2010 Doran ..... H04L 41/12  
370/254  
2013/0050246 A1\* 2/2013 Barry ..... G06F 11/32  
345/593  
2014/0013432 A1 1/2014 Yi  
2014/0181718 A1\* 6/2014 Gao ..... G06F 3/0484  
715/771  
2014/0201642 A1\* 7/2014 Vicat-Blanc ..... H04L 41/22  
715/736

2015/0207755 A1\* 7/2015 Tsai ..... H04L 47/76  
709/226  
2015/0277982 A1\* 10/2015 Iwata ..... G06F 9/505  
718/103  
2016/0119204 A1\* 4/2016 Murasato ..... G06F 3/0481  
715/735  
2016/0212172 A1\* 7/2016 Senanayake ..... H04L 63/20  
2016/0219078 A1\* 7/2016 Porras ..... G06F 3/0484  
2017/0034274 A1\* 2/2017 Hanney ..... H04L 67/1097

OTHER PUBLICATIONS

Estrin et al., "Network Visualization with Nam, the VINT Network Animator," Computer, Nov. 2000, pp. 63-68, IEEE.  
Chopade et al., "Visualization Techniques for Large-Scale Big Data Networks: Smart Power Grid Survivability in a Complex Operating Environment," ASE BigData/SocialInformatics/PASSAT/BioMedCom 2014 Conference, Dec. 14-16, 2014, ASE.  
Shi et al., "1.5D Egocentric Dynamic Network Visualization," IEEE Transactions on Visualization and Computer Graphics, May 2015, pp. 624-637, vol. 21, No. 5.

\* cited by examiner

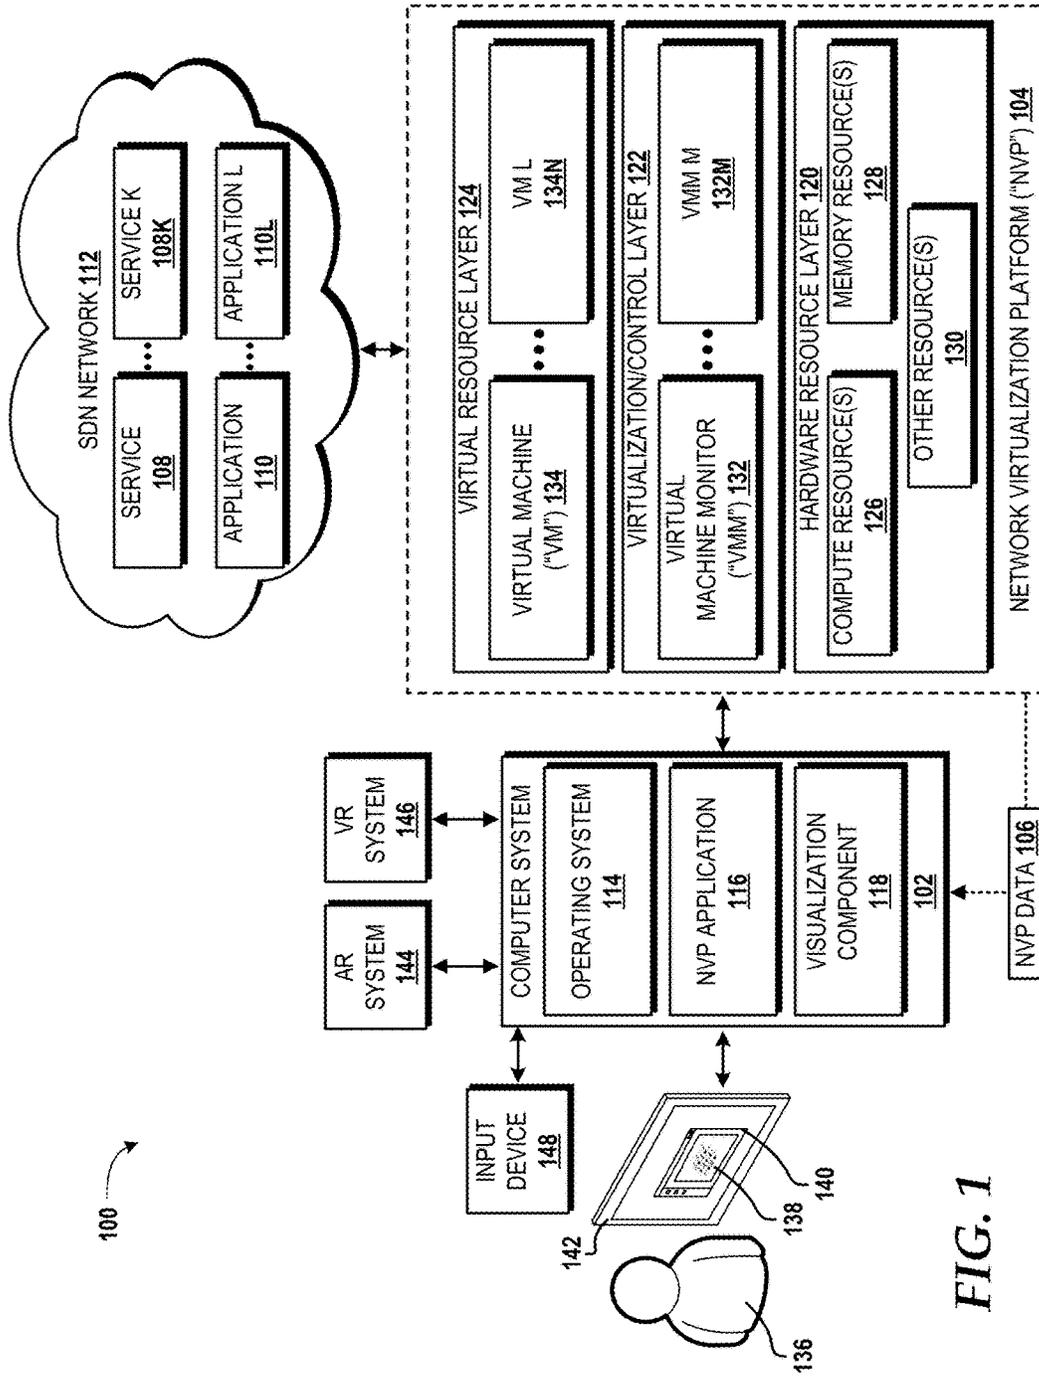


FIG. 1

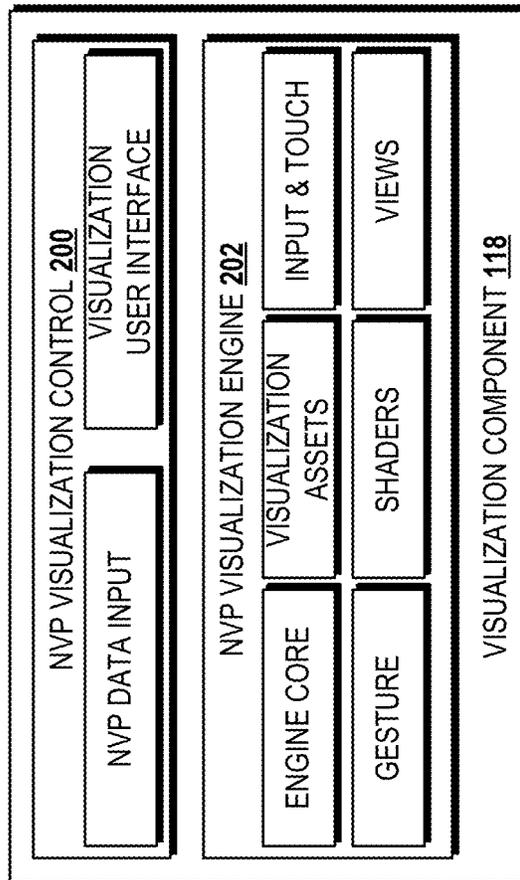
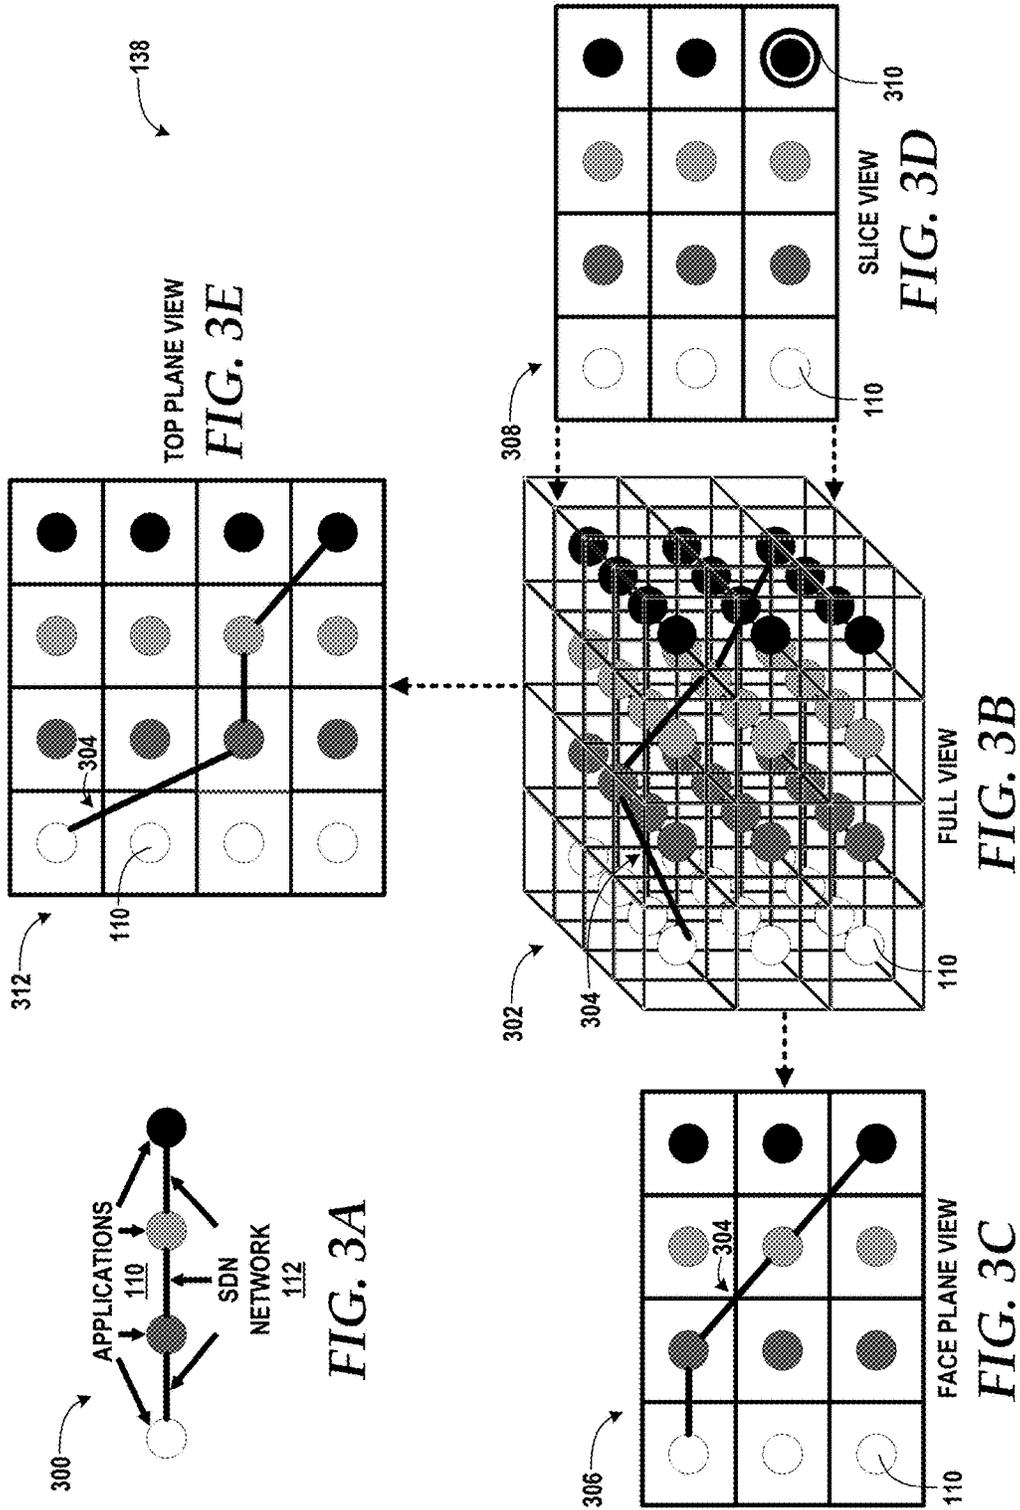
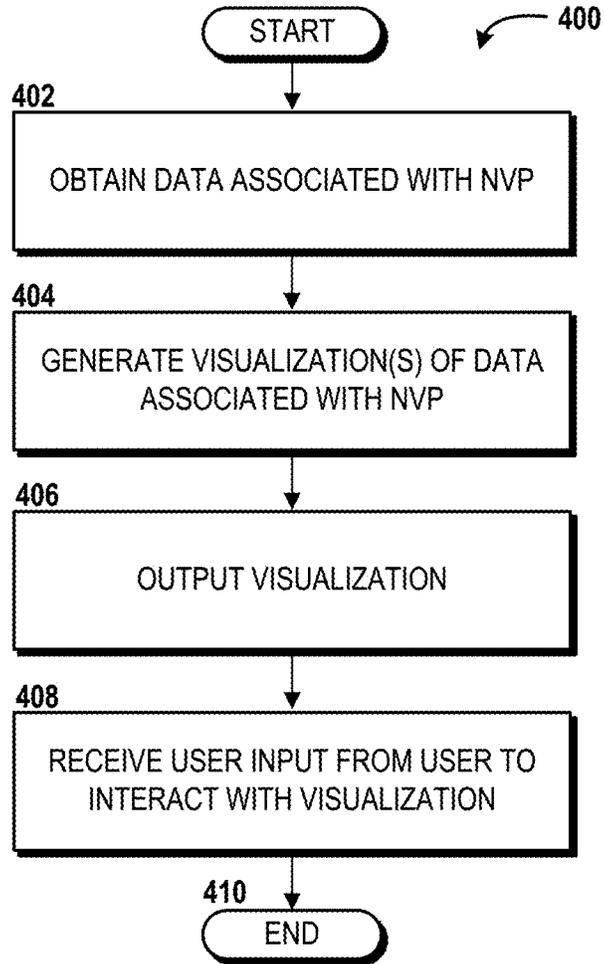
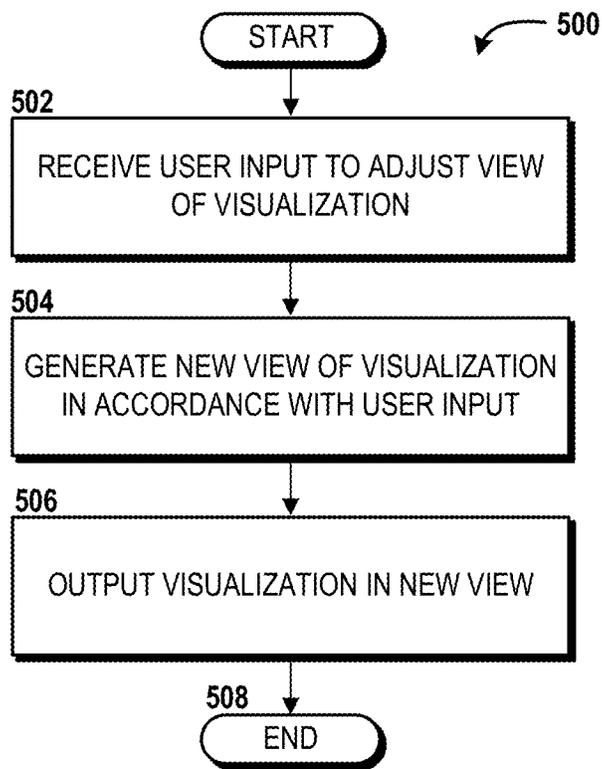


FIG. 2

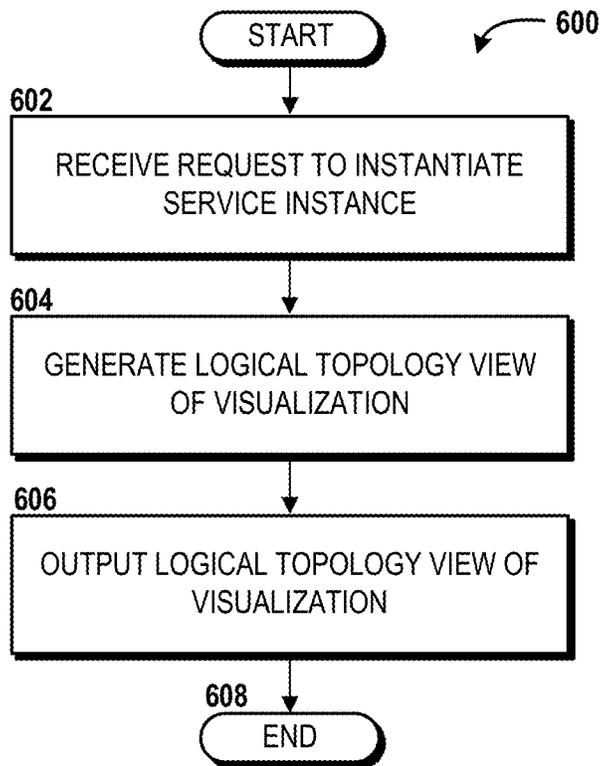




**FIG. 4**



*FIG. 5*



**FIG. 6**

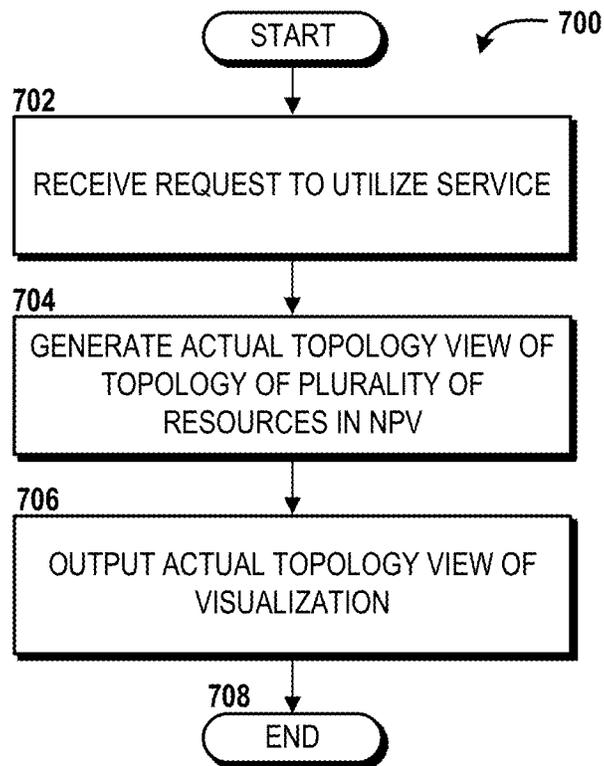


FIG. 7

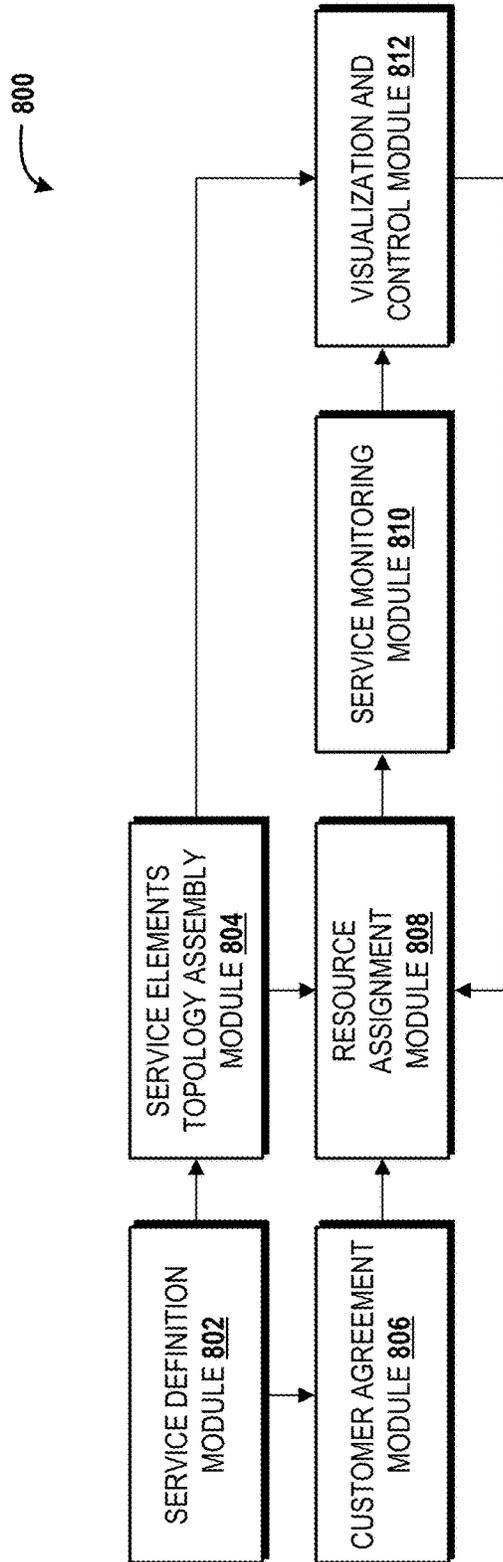


FIG. 8

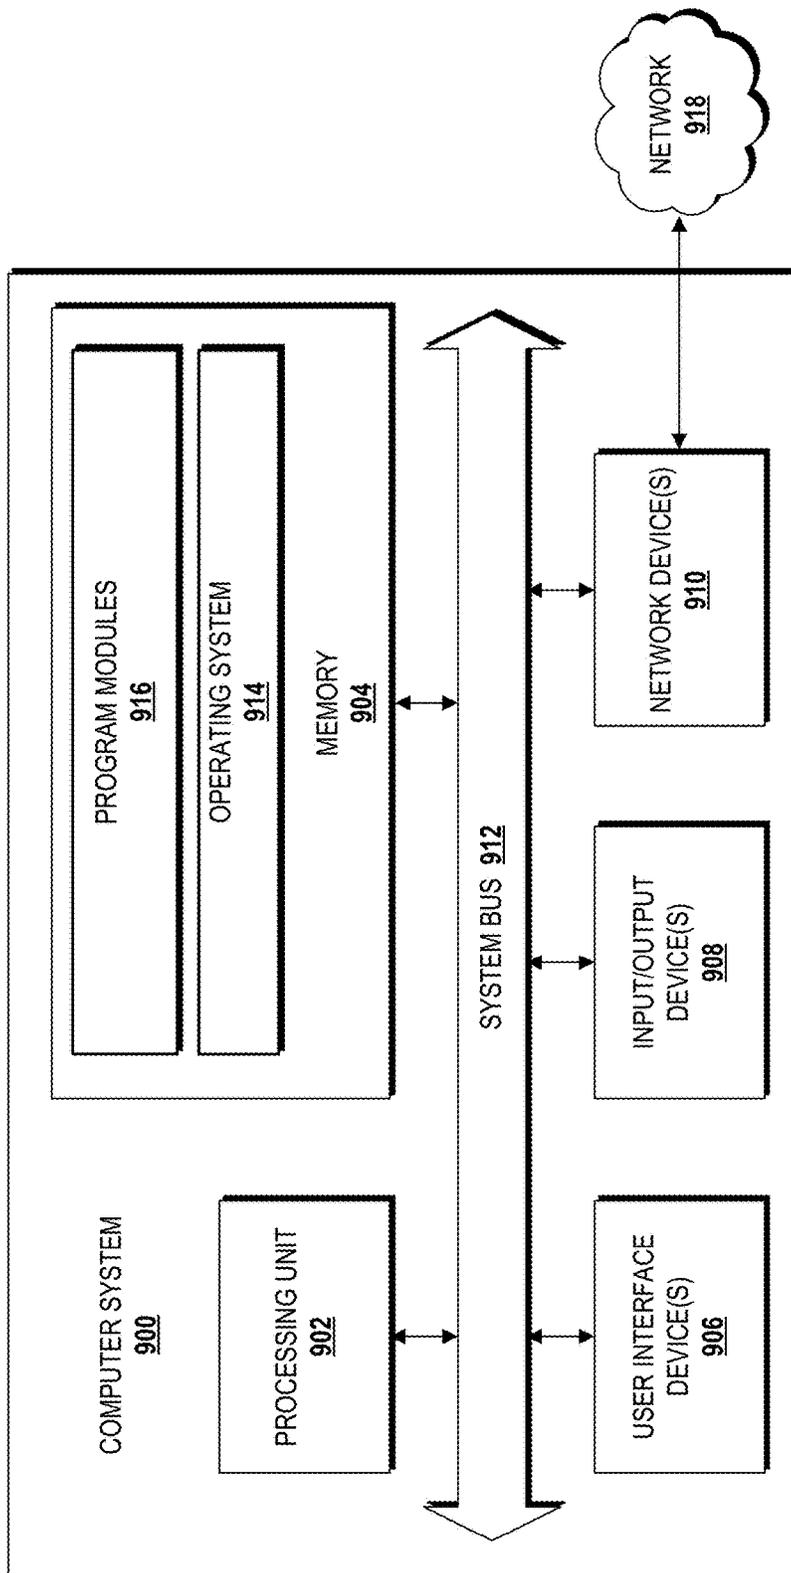


FIG. 9

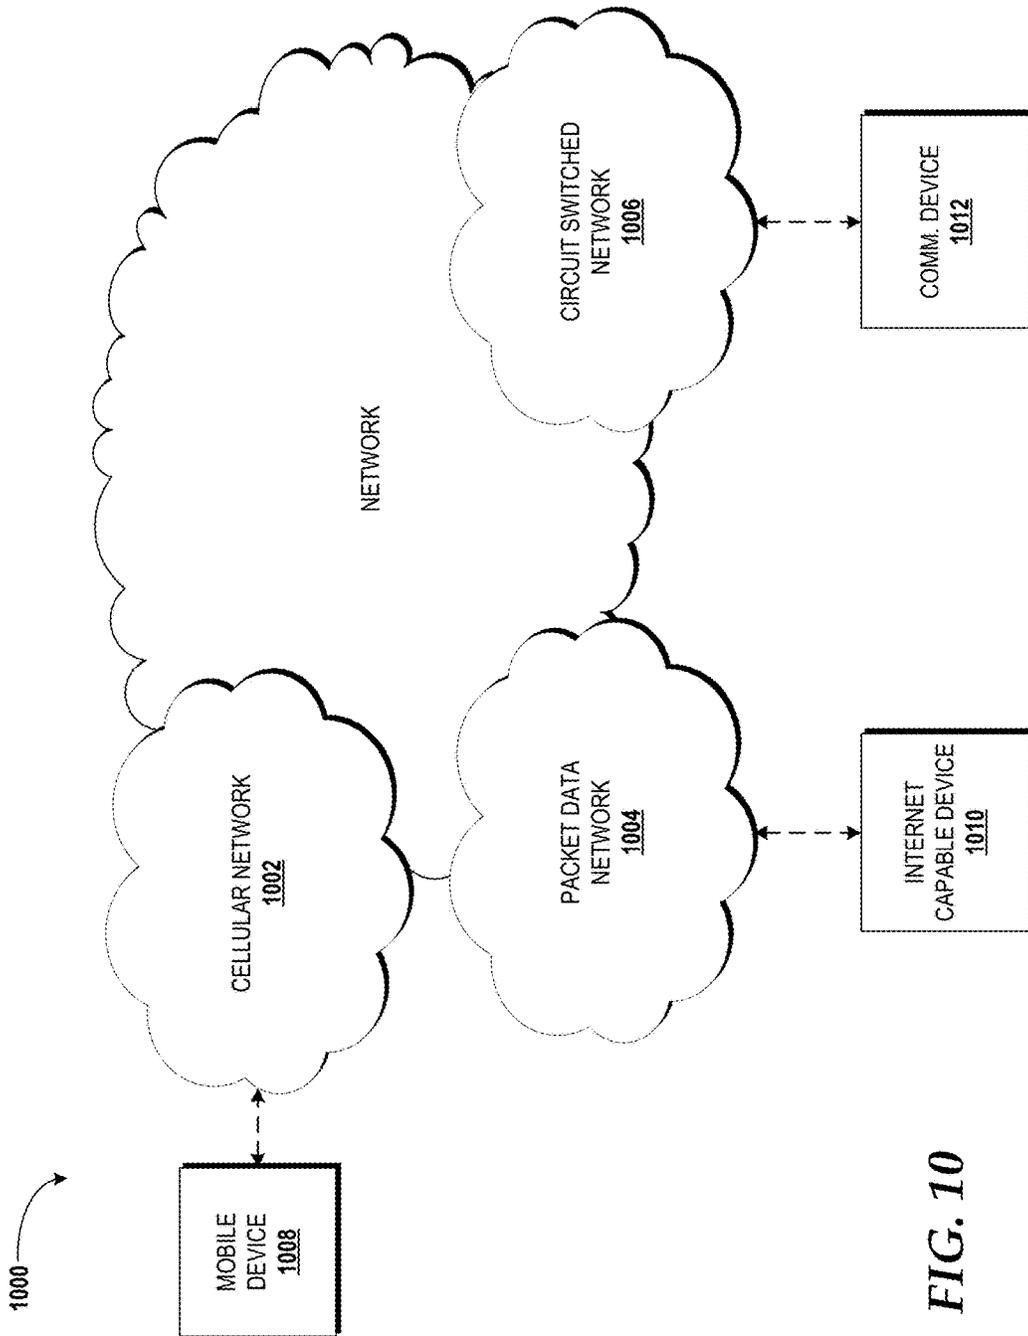


FIG. 10

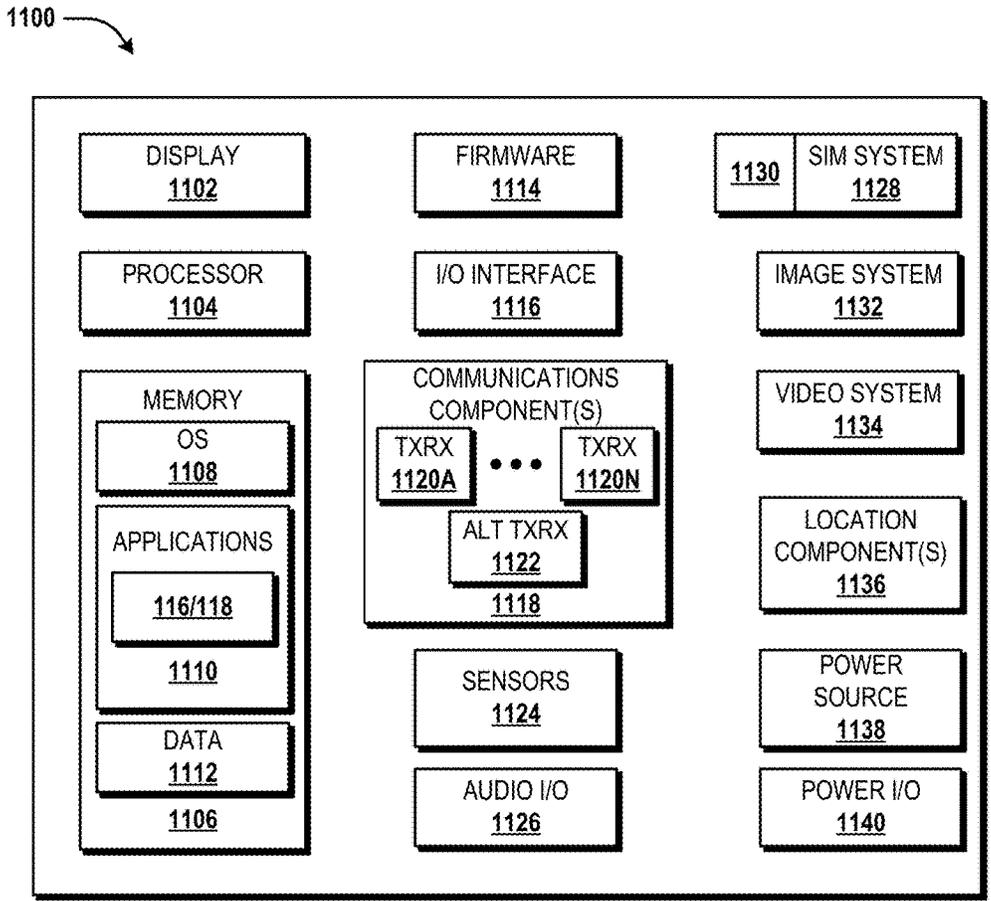


FIG. 11

## VISUALIZATION FOR NETWORK VIRTUALIZATION PLATFORM

### BACKGROUND

Software-defined networking (“SDN”) is an architectural framework for creating intelligent networks that are programmable, application aware, and more open. SDN provides an agile and cost-effective communications platform for handling the dramatic increase in data traffic on carrier networks by providing a high degree of scalability, security, and flexibility. SDN provides several benefits. SDN can allow for the creation of multiple virtual network control planes on common hardware. SDN can help extend service virtualization and software control into many existing network elements. SDN enables applications to request and manipulate services provided by the network and to allow the network to expose network states back to the applications. SDN exposes network capabilities through application programming interfaces (“APIs”), making the control of network equipment remotely accessible and modifiable via third-party software clients using open protocols such as OpenFlow, available from Open Network Forum (“ONF”).

User-defined, on-demand cloud services and user digital experience expectations are driving planning and deployment of network functional virtualization and service-centric SDN among global telecommunications service providers. Network Virtualization Platforms (“NVPs”) are deployed in information technology (“IT”) data centers, network central offices, and other network points of presence (“POPs”) to accelerate deployment of on-demand user service and virtualized network functions. An NVP is a shared virtualized infrastructure that supports multiple services and network applications (including real-time and non-real-time applications).

Combining SDN and NVP functionality, such as in Domain 2.0, available from AT&T, provides a highly complex and dynamic set of relationships between virtual, logical, and physical resources. Autonomous controls and real-time service path steering, workload, creation, distribution, and destruction are inherent Domain 2.0 capabilities, but currently no unified and holistic method exists for human operators to understand, model, and design service paths and to understand and control the operations of the environment. Currently available solutions are dominated by tools and techniques developed to support traditional information technology and network domains, which, in comparison, are relatively simple and static, and are therefore ill-suited to Domain 2.0 and similar environments. The new dimensions of Domain 2.0 and similar environments require a new dimension of human insight and control.

### SUMMARY

Concepts and technologies disclosed herein are directed to visualizations for network virtualization platforms (“NVPs”). According to one aspect disclosed herein, a computer system executing a visualization component can obtain data associated with an NVP. The computer system can generate a visualization of the data associated with the NVP. The computer system can output the visualization.

In some embodiments, the visualization includes a three-dimensional visualization. The three-dimensional visualization can include a central cube visualization. The central cube visualization can include a plurality of smaller cube visualizations. Each of the plurality of smaller cube visualizations can include a node representative of an application.

The application can be executable by at least a portion of a plurality of resources of the NVP. The central cube visualization, in some embodiments, includes a plurality of networking links between at least a portion of the plurality of applications of the plurality of smaller cube visualizations.

In some embodiments, the data includes data associated with a topology of the plurality of resources of the NVP. In these embodiments, the computer system can receive a request to instantiate a service instance of a service provided by the NVP. In response to the request, the computer system can generate a logical topology view of the central cube visualization. The computer system can output the logical topology view of the central cube visualization. The logical topology view can include at least a portion of the topology of the plurality of resources of the NVP. The logical topology view can represent a service path. The service path can include and identify a set of the plurality of resources utilized by the service.

The logical topology view, in some embodiments, can include a face view of the central cube visualization. The face view can include a plurality of layers. Each layer of the plurality of layers can represent a physical technical plant location and a subset of the set of the plurality of resources utilized by the service set of the plurality of resources utilized by the service and located within the physical technical plant (“TeP”) location. Physical TePs can be or can include one or more telco central offices (“COs”), one or more data centers, one or more outside plant environments (e.g., huts, bunkers, poles, and the like), customer premises locations, and the like.

In some embodiments, the computer system can receive a request to utilize a service. The computer system can generate an actual topology view of at least a portion of the topology of the plurality of resources of the NVP. The actual topology view can represent a service graph. The service graph can include a set of the plurality of resources utilized by the service. The computer service can output the actual topology view.

It should be appreciated that the above-described subject matter may be implemented as a computer-controlled apparatus, a computer process, a computing system, or as an article of manufacture such as a computer-readable storage medium. These and various other features will be apparent from a reading of the following Detailed Description and a review of the associated drawings.

This Summary is provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended that this Summary be used to limit the scope of the claimed subject matter. Furthermore, the claimed subject matter is not limited to implementations that solve any or all disadvantages noted in any part of this disclosure.

### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram illustrating aspects of an illustrative operating environment for implementing the various concepts and technologies disclosed herein.

FIG. 2 is a block diagram illustrating aspects of a visualization component used by a computer system to generate a visualization of data associated with an NVP, according to an illustrative embodiment.

FIGS. 3A-3E are diagrams illustrating aspects of various views of an example visualization, according to an illustrative embodiment.

FIG. 4 is a flow diagram illustrating aspects of a method for generating a visualization of NVP data, according to an illustrative embodiment.

FIG. 5 is a flow diagram illustrating aspects of a method for adjusting a new view of a visualization of NVP data, according to an illustrative embodiment.

FIG. 6 is a flow diagram illustrating aspects of a method for generating a logical topology view of a visualization of NVP data, according to an illustrative embodiment.

FIG. 7 is a flow diagram illustrating aspects of a method for generating an actual topology view of a visualization of NVP data, according to an illustrative embodiment.

FIG. 8 is a block diagram illustrating a plurality of modules associated with an NVP application and a visualization component that can be executed by a computer system to perform various operations to facilitate user interaction with a visualization, according to an illustrative embodiment.

FIG. 9 is a block diagram illustrating an example computer system capable of implementing aspects of the embodiments presented herein.

FIG. 10 schematically illustrates a network, according to an illustrative embodiment.

FIG. 11 is a block diagram illustrating an example mobile device and components thereof, according to an illustrative embodiment.

#### DETAILED DESCRIPTION

Combining SDN and NVP functionality, such as in Domain 2.0, available from AT&T, provides a highly complex and dynamic set of relationships between virtual, logical, and physical resources. Autonomous controls and real-time service path steering, workload, creation, distribution, and destruction are inherent Domain 2.0 capabilities but currently no unified and holistic method exists for human operators to understand, model, and design service paths and to understand and control the operations of the environment. Currently available solutions are dominated by tools and techniques developed to support traditional information technology and network domains, which, in comparison, are relatively simple and static, and are therefore ill-suited to Domain 2.0 and similar environments. The new dimensions of Domain 2.0 and similar environments require a new dimension of human insight and control.

The concepts and technologies facilitate relating physical and logical aspects of Domain 2.0 and similar networks and applications for use in service design, instantiation, and operations by both internal users and external customers. The concepts and technologies described herein, in some embodiments, provide a three-dimensional cube metaphor that allows visualization, navigation, and manipulation of service logic specified by the alignment of service paths onto a service graph. Moreover, the application of virtual reality and augmented reality technologies provides an intuitive, effective, and highly-differentiated experience for network management.

While the subject matter described herein may be presented, at times, in the general context of program modules that execute in conjunction with the execution of an operating system and application programs on a computer system, those skilled in the art will recognize that other implementations may be performed in combination with other types of program modules. Generally, program modules include routines, programs, components, data structures, computer-executable instructions, and/or other types of structures that perform particular tasks or implement par-

ticular abstract data types. Moreover, those skilled in the art will appreciate that the subject matter described herein may be practiced with other computer systems, including handheld devices, mobile devices, wireless devices, multiprocessor systems, distributed computing systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers, routers, switches, other computing devices described herein, and the like.

While connections are shown between some of the components illustrated in FIG. 1, it should be understood that some, none, or all of the components illustrated in FIG. 1 can be configured to interact with one other to carry out various functions described herein. Thus, it should be understood that FIG. 1 and the following description are intended to provide a general understanding of a suitable environment in which various aspects of embodiments can be implemented, and should not be construed as being limiting in any way.

Referring now to FIG. 1, aspects of an operating environment 100 for implementing various embodiments of the concepts and technologies disclosed herein will be described. The operating environment 100 shown in FIG. 1 includes a computer system 102 operating on or in communication with an NVP 104 to provide functionality in accordance with the concepts and technologies disclosed herein to visualize NVP data 106 associated with and received from the NVP 104, including, for example, data associated with one or more services 108-108K provided, at least in part, by one or more applications 110-110L operating within an SDN network 112 created and managed, at least in part, by the NVP 104.

According to various implementations of the concepts and technologies disclosed herein, the computer system 102 can include a user computing device, such as a tablet computing device, a personal computer ("PC"), a desktop computer, a laptop computer, a notebook computer, a cellular phone or smartphone, other mobile computing devices, a personal digital assistant ("PDA"), or the like. Example architectures of the computer system 102 are illustrated and described herein below with reference to FIGS. 8 and 10. The functionality of the computer system 102 can be provided, at least in part, by a cloud-based computing platform that can be provided by one or more application servers, web servers, data storage systems, network appliances, dedicated hardware devices, and/or other server computers or computing devices, which might, in some embodiments, leverage compute, memory, and/or other resources of the NVP 104 to perform operations described herein. In light of the above alternative embodiments of the computer system 102, it should be understood that this example is illustrative and should not be construed as being limiting in any way.

The computer system 102 can be configured to execute an operating system 114 and one or more application programs such as, for example, an NVP application 116, a visualization component 118, and/or other application programs. The operating system 114 is a computer program for controlling the operation of the computer system 102. The application programs are executable programs configured to execute on top of the operating system 114. In some embodiments, at least a portion of the NVP application 116 functionality and/or the visualization component 118 functionality is included in the operating system 114. In some other embodiments, at least a portion of the NVP application 116 functionality and/or the visualization component 118 functionality is included in the NVP 104, wherein the computer system 102 accesses the functionality via a communications interface.

In some embodiments, the visualization component **118** can be implemented as a component of the NVP application **116**, and in some embodiments, the visualization component **118** can be implemented as a component separate from the NVP application **116**. In some embodiments, the visualization component **118** can be provided as a stand-alone application. Thus, while the NVP application **116** and the visualization component **118** are illustrated as components of the computer system **102**, it should be understood that each of these components, or combinations thereof, may be embodied as or in stand-alone devices or components thereof operating on or in communication the computer system **102**. Thus, the illustrated embodiment is illustrative, and should not be construed as being limiting in any way.

In some embodiments, the visualization component **118** may be implemented as a plugin or add-in for the NVP application **116**. In some other embodiments, the visualization component **118** can include a service and/or set of application programming interfaces (“APIs”) that can provide the functionality described herein. Thus, it should be appreciated that the visualization component **118** can be implemented as hardware, software, or a combination thereof.

The NVP **104** is a shared infrastructure that can support multiple services, such as the services **108**, and multiple network applications, such as the applications **110** (including real-time and non-real-time applications). The NVP **104** can utilize cloud sharing constructs to virtualize network functions by decoupling hardware and software, making networks, such as the SDN network **112**, more flexible and physically uniform by minimizing dependence upon hardware constraints. The NVP **104** uses SDN to provide programmability and abstraction of underlying hardware complexity, and to separate management and control planes from the data plane.

The applications **110** can include, but are not limited to, one or more applications that provide or utilize, at least in part, one or more of the services **108**, such as, for example, domain name service (“DNS”), remote authentication dial-in user service (“RADIUS”), dynamic host configuration protocol (“DHCP”), lightweight directory access protocol (“LDAP”), content routing in either or both of the control and data planes, and route reflecting in either or both the control and data planes. The services **108** also can include one or more real-time services such as, but are not limited to, voice over internet protocol (“VoIP”) service, streaming video service, videoconferencing service, online gaming service, chatting service, instant messaging (“IM”) service, and the like in the service plane. Both synchronous and asynchronous (e.g., transactional) services are contemplated.

The illustrated NVP **104** includes a hardware resource layer **120**, a virtualization/control layer **122**, and a virtual resource layer **124** that work together to perform operations as will be described in detail herein. The hardware resource layer **120** provides hardware resources, which, in the illustrated embodiment, include one or more compute resources **126**, one or more memory resources **128**, and one or more other resources **130**.

The compute resource(s) **126** can include one or more hardware components that perform computations to process data, and/or to execute computer-executable instructions of one or more application programs, operating systems, and/or other software, including the applications **110** to provide, at least in part, the services **108**. The compute resources **126** can include one or more central processing units (“CPUs”) configured with one or more processing cores. The compute

resources **126** can include one or more graphics processing unit (“GPU”) configured to accelerate operations performed by one or more CPUs, and/or to perform computations to process data, and/or to execute computer-executable instructions of one or more application programs, operating systems, and/or other software that may or may not include instructions particular to graphics computations. In some embodiments, the compute resources **126** can include one or more discrete GPUs. In some other embodiments, the compute resources **126** can include CPU and GPU components that are configured in accordance with a co-processing CPU/GPU computing model, wherein the sequential part of an application executes on the CPU and the computationally-intensive part is accelerated by the GPU. The compute resources **126** can include one or more system-on-chip (“SoC”) components along with one or more other components, including, for example, one or more of the memory resources **128**, and/or one or more of the other resources **130**. In some embodiments, the compute resources **126** can be or can include one or more SNAPDRAGON SoCs, available from QUALCOMM of San Diego, Calif.; one or more TEGRA SoCs, available from NVIDIA of Santa Clara, Calif.; one or more HUMMINGBIRD SoCs, available from SAMSUNG of Seoul, South Korea; one or more Open Multimedia Application Platform (“OMAP”) SoCs, available from TEXAS INSTRUMENTS of Dallas, Tex.; one or more customized versions of any of the above SoCs; and/or one or more proprietary SoCs. The compute resources **126** can be or can include one or more hardware components architected in accordance with an ARM architecture, available for license from ARM HOLDINGS of Cambridge, United Kingdom. Alternatively, the compute resources **126** can be or can include one or more hardware components architected in accordance with an x86 architecture, such an architecture available from INTEL CORPORATION of Mountain View, Calif., and others. Those skilled in the art will appreciate the implementation of the compute resources **126** can utilize various computation architectures or combinations thereof, and as such, the compute resources **126** should not be construed as being limited to any particular computation architecture or combination of computation architectures, including those explicitly disclosed herein.

The memory resource(s) **128** can include one or more hardware components that perform storage operations, including temporary or permanent storage operations. In some embodiments, the memory resource(s) **128** include volatile and/or non-volatile memory implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules, or other data disclosed herein. Computer storage media includes, but is not limited to, random access memory (“RAM”), read-only memory (“ROM”), erasable programmable ROM (“EPROM”), electrically erasable programmable ROM (“EEPROM”), flash memory or other solid state memory technology, CD-ROM, digital versatile disks (“DVD”), or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store data and which can be accessed by the compute resources **126**.

The other resource(s) **130** can include any other hardware resources that can be utilized by the compute resources(s) **126** and/or the memory resource(s) **128** to perform operations described herein. The other resource(s) **130** can include one or more input and/or output processors (e.g., network interface controller or wireless radio), one or more modems, one or more codec chipset, one or more pipeline processors,

one or more fast Fourier transform (“FFT”) processors, one or more digital signal processors (“DSPs”), one or more speech synthesizers, and/or the like.

The hardware resources operating within the hardware resource layer **120** can be virtualized by one or more virtual machine monitors (“VMMs”) **132-132M** (also known as “hypervisors;” hereinafter “VMMs **132**”) operating within the virtualization/control layer **122** to manage one or more virtual resources that reside in the virtual resource layer **124**. The VMMs **132** can be or can include software, firmware, and/or hardware that alone or in combination with other software, firmware, and/or hardware, manages one or more virtual resources operating within the virtual resource layer **124**.

The virtual resources operating within the virtual resource layer **124** can include abstractions of at least a portion of the compute resources **126**, the memory resources **128**, the other resources **130**, or any combination thereof. These abstractions are referred to herein as virtual machines (“VMs”). It should be understood, however, that other container technologies can be used and are contemplated. In the illustrated embodiment, the virtual resource layer **124** includes VMs **134-134N** (hereinafter “VMs **134**”). The VMs **134** can execute the applications **110** to provide, at least in part, the services **108**. Each of the VMs **134** can execute one or more of the applications **110** or one or more portions thereof. For example, one or more of the VMs **134** can execute one or more of the applications **110** to utilize or provide, at least in part, the services **108** such as, but not limited to, DNS, RADIUS, DHCP, LDAP, content routing in either or both of the control and data planes, and route reflecting in either or both the control and data planes. The applications **110** also can include one or more applications that provide, at least in part, one or more real-time services such as, but are not limited to, VoIP, streaming video service, videoconferencing service, online gaming service, chatting service, IM service, and the like in the service plane.

The NVP data **106** can include data any data associated with the NVP **104**. For example, the NVP data **106** can be or can include data associated with the operational state(s) and configuration(s) of one or more components of the NVP **104**, such as the operational states and configurations of the virtual resources operating within the virtual resource layer **124**. The NVP data **106** also can include data associated with application resources, such as one or more of the applications **110-110L**.

The NVP **104** can provide the NVP data **106** to the computer system **102** periodically, continuously, and/or upon request. The NVP application **116** can, in some embodiments, include functionality to enable a user **136** to schedule periodic retrieval of the NVP data **106** from the NVP **104** or to request the NVP data **106** on-demand. The NVP **104**, in some embodiments, can be configured to report the NVP data **106** to the computer system **102** periodically or in response to an event, such as a change of an operational state and/or of a configuration of one or more of the virtual resources operating within the virtual resource layer **124** and/or one or more application resources instantiated in the SDN network **112**.

The NVP application **116** can receive the NVP data **106** and can provide at least a portion of the NVP data **106** to the visualization component **118**. The visualization component **118** alternatively can request the NVP data **106** from the NVP application **116**. In either case, the visualization component **118** can obtain the NVP data **106**, generate, based upon the NVP data **106**, a visualization **138** of the NVP data **106**, and output the visualization **138** via a user interface

(“UI”) **140** of the visualization component **118** via a display **142**, an augmented reality (“AR”) system **144**, a virtual reality (“VR”) system **146**, or a combination thereof. The user **136** can interact with the UI **140** of the visualization component **118** to view and interact with the visualization **138**. In doing so, the user **136** can exert influence upon the automation and control of the NVP **104** or at least a portion thereof.

The visualization **138** can be generated by any software framework designed for the creation and development of graphics. Some example software frameworks include, but are not limited to, AUTOCAD (available from AUTODESK), UNREAL ENGINE (available from EPIC GAMES, INC.), UNITY (available from UNITY TECHNOLOGIES), CRYENGINE (available from CRYTEK), HAVOK VISION ENGINE (available from HAVOK), other proprietary software frameworks, open source software frameworks, combinations thereof, and the like. Those skilled in the art will appreciate the wide range of graphical fidelity, visual styles, and other attributes the virtualization **138** might have, including color, shape, texture, space, form, dynamics (e.g., bouncing, flashing, and the like), and other design and/or interaction elements. As such, further details in this regard are not provided herein since these are mere design decisions.

The display **142** is an output device configured to present information in a visual form. In some embodiments, the display **142** is a liquid crystal display (“LCD”) utilizing any active or passive matrix technology and can utilize any backlighting technology. In some embodiments, the display **142** is an organic light emitting diode (“OLED”) display. The display **142** can be embodied with other display technologies. As such, the examples provided above should not be considered limiting in any way.

The AR system **144** can provide an augmented reality through which at least a portion of a physical, real-world environment is augmented to include the visualization **138**. The visualization **138** can be presented via the AR system **144** over and/or spatially integrated with real-world objects of the physical, real-world environment. In these embodiments, the AR system **144** can utilize a camera component (not shown) to provide a live view of the physical, real-world environment to be augmented with the visualization **138**. In other embodiments, the AR system **144** can provide a non-live view of a physical, real-world environment. The non-live view can present a physical, real-world environment as a static image representative of a past reality that can be augmented with the visualization **138**.

The AR system **144** can be in communication with the computer system **102** via a wireless or wired connection through which data, such as the NVP data **106** and other data associated with the visualization **138**, can be shared. The AR system **144** can function as a stand-alone system that utilizes on-board computing components to perform operations to present and facilitate manipulation of an augmented reality augmented with the visualization **138**, or alternatively, can leverage the computing resources of the computer system **102**, or even at least a portion of the hardware resources in the hardware resource layer **120** of the NVP **104** to perform such operations.

The AR system **144**, in some embodiments, is or includes a camera (e.g., a still camera and/or video camera), a sensor (e.g., an accelerometer, a global positioning system sensor, a solid state compass, or the like), a display (e.g., an integrated display, a head-mounted display, an eyeglasses display, a head-up display, an external monitor, a projection system, or a holographic display), an input device, the like,

or any combination thereof. In some embodiments, the AR system 144 is or includes a wearable computing device that includes an integrated display through which to present the visualization 138 via an augmented reality. The AR system 144, in these embodiments, can be GOOGLE GLASS, available from GOOGLE INC., or MICROSOFT HOLOLENS, available from MICROSOFT CORP. Other devices, such as mobile telephones, smartphones, tablet computers, slate computers, smart watches, laptop computers, notebook computers, ultrabook computers, netbook computers, computers of other form factors, computing devices of other form factors, other computing systems, other computing devices, and/or the like, can be configured to execute one or more AR applications to function as the AR system 144 described herein.

The VR system 146 can provide a computer-generated environment (also referred to herein as a “virtual” environment) that the user 136 can explore. The virtual environment can be or can include the visualization 138. A VR environment can include a computer-generated representation or at least an approximation of at least a portion of a physical, real-world environment with the visualization 138 integrated within. The VR environment can be at least partially different from the physical, real-world environment of which the VR environment is representative with the visualization 138 integrated within. The VR environment can include virtual objects not found in the corresponding physical, real-world environment. Lighting effects such as light bloom and other effects such as depth-of-field can be applied to the VR environment to create atmosphere. Moreover, natural phenomena such as gravity and momentum can be simulated in the virtual environment. These natural phenomena can be simulated, for example, when the user 136 interacts with the visualization 138.

The VR system 146, in some embodiments, is or includes a display (e.g., an integrated display, a head-mounted display, an eyeglasses display, a head-up display, an external monitor, or a projection system), an input device, a combination thereof, or the like. In some embodiments, the VR system 146 is OCULUS RIFT (available from OCULUS VR), SAMSUNG GEAR VR (available from SAMSUNG and OCULUS VR), GOOGLE CARDBOARD (available from GOOGLE), HTC VIVE (available from HTC), MORPHEUS (available from SONY), or the like.

The visualization 138 can be presented on the display 142, via the AR system 144, and/or the VR system 146 and can be interacted with by the user 136 through one or more input devices 148. Although shown as being external to the computer system 102, the display 142, the AR system 144, and the VR system 146, the input device(s) 148 can be implemented as part of the computer system 102, the display 142, the AR system 144, and/or the VR system 146. The input device(s) 148 can be or can include one or more touchscreens, one or more multi-touch touchscreens, one or more keyboards, one or more computer mice, one or more game controllers, one or more joysticks, one or more touchpads, one or more gesture devices (e.g., MICROSOFT KINECT, available from MICROSOFT CORPORATION), combinations thereof, or the like.

The visualization 138, in some embodiments, includes a three-dimensional visualization. The three-dimensional visualization can include a central cube visualization. The central cube visualization can include a three-dimensional assemblage of stacks of smaller cubes, each representing one or more application resources such as one or more of the applications 110. At the core of each smaller cube in the visualization 138 can be a node representing the application

resource(s). Networking among the application resources can be represented as links between the nodes of the cubes. When referencing the visualization 138 in the description below, a central cube visualization as described above will provide the context. It should be understood, however, that other shapes can be used by the visualization component 118 to convey, to the user 136, similar information. Some illustrative examples of the visualization 138 are illustrated with reference to FIG. 3, which is described in greater detail herein below.

The visualization 138 can provide a service graph. A “service graph,” as used herein, is a collection of available resources (such as application resources embodied as the applications 110) along with sequencing and relationships available to be used to process data plane service requests to instantiate a service instance of one or more of the services 108. A service graph describes the actual topology of the resources instantiated and operationalized by the NVP 104 within the SDN network 112. A service graph can be viewable as a full view of the visualization 138. An illustrative example of a service graph is shown in FIG. 3B.

A “service path,” as used herein, is a collection of specific resources from the service graph along with sequencing and relationships used to process a specific data plane service request. A service path describes a logical topology of the resources included in a particular instance of a service, such as one of the services 108. The illustrative service paths provided herein follow a flow convention of left-to-right, although right-to-left and/or other flows are implemented. A logical view of a service path can be provided by a face plane view of the visualization 138. An illustrative example of a logical view of a service path as a face plane view of a central cube visualization is shown in FIG. 3C. The visualization 138 can provide a plurality of layers of the face plane view. Each of the layers can represent a physical technical plant (“TeP”) location. Physical TePs can be or can include one or more telco central offices (“COs”), one or more data centers, one or more outside plant environments (e.g., huts, bunkers, poles, and the like), customer premises locations, and the like. The specific resources in a specific TeP location that are processing a particular data plane service request can be identified by sliding out the corresponding layer for examination by the user 136, or by navigating into the visualization 138.

A network (or administrative domain) wide view across all TePs can be provided by a top plane view of the visualization 138. A top plane view identifies the collection of all potential physical resources, such as provided in the hardware resource layer 120 of the NVP 104, available to process a particular data plane service request, as well as the specific resources that are processing a particular data plane service request. This view is perhaps most valuable for steering service requests, which would be enabled by disconnecting a network link from an application resource in one TeP and reconnecting it to the same type of resource in another TeP. An illustrative example of this view is shown in FIG. 3E. Service logic can be defined by aligning a service path onto a service graph. “Service Logic,” as used herein, is what aligns a service path onto a service graph. Data plane traffic steering/engineering is one aspect of service logic and can be accomplished by adding and connecting application resources across and within TePs. Service logic can be constrained by affinity/anti-affinity rules, quality of service (“QoS”) parameters (e.g., latency), regulatory, business needs, and the like. Service logic controls embedded within the visualization 138 can be manipulated by operators and/or

customers, such as the user 136, to influence the data plane traffic of the underlying NVP, such as the NVP 104.

The user 136 can interact with the visualization 138 through the UI 140 using various actions. The actions can include user input received by the computer system 102 via the display 142, the AR system 144, the VR system 146, the input device(s) 148, or some combination thereof from the user 136. The actions can include movement actions such as, but not limited to, pan up, pan down, pan left, pan right, rotate clockwise, rotate counter-clockwise, tilt down, tilt up, zoom in, zoom out, field-of-view increase, field-of-view decrease, and the like. The movement actions can be used to navigate the visualization 138. Other actions to add, delete, and/or modify at least a portion of the visualization 138 also are contemplated.

In some embodiments, the user 136 might perform gestures as user input to perform one or more actions. Gestures might be performed on a single or multi-touch touchscreen. Gestures might be open-space gestures through which the user 136 utilizes at least a portion of his or her body as user input to perform one or more actions in space. For example, the input device 148 embodied as a KINECT device can enable the user 136 to perform such gestures to interact with the visualization 138. Several gestures will now be described. It should be understood that these gestures are illustrative and are not intended to limit the scope of the appended claims. Moreover, the described gestures, additional gestures, and/or alternative gestures may be implemented in software for use by the computer system 102 via the display 142, the AR system 144, the VR system 146, the input device(s) 148, or some combination thereof. As such, a developer may create gestures that are specific to a particular application program, such as the visualization component 118.

In some embodiments, the input device 148 embodied as a touchscreen supports a tap gesture in which the user 136 taps the touchscreen once on an item presented on the display 142 to select at least a portion of the visualization 138. In some embodiments, the touchscreen supports a double tap gesture in which the user 136 taps the touchscreen twice on an item presented on the display 142 to perform an action such as, but not limited to, zooming into or zooming out of the visualization 138 in stages. In some embodiments, the touchscreen supports a tap and hold gesture in which the user 136 taps the touchscreen and maintains contact for at least a pre-defined time to perform an action such as opening a context-specific menu to make available additional options for the user 136 to interact with the virtualization 138.

In some embodiments, the input device 148 embodied as a touchscreen supports a pan gesture in which the user 136 places a finger on the touchscreen and maintains contact with the touchscreen while moving the finger on the touchscreen. The pan gesture may be used for various reasons including, but not limited to, moving through or around the visualization 138 at a controlled rate. Multiple finger pan gestures are also contemplated. In some embodiments, the touchscreen supports a flick gesture in which the user 136 swipes a finger in the direction the user 136 wants the screen to move. The flick gesture may be used for various reasons including, but not limited to, scrolling horizontally or vertically through the visualization 138. In some embodiments, the touchscreen supports a pinch and stretch gesture in which the user 136 makes a pinching motion with two fingers (e.g., thumb and forefinger) on the touchscreen or moves the two fingers apart. The pinch and stretch gesture

may be used for various reasons including, but not limited to, zooming gradually in or out of the visualization 138.

Although the above gestures have been described with reference to the use of one or more fingers for performing the gestures on the input device 148 embodied as a touchscreen, other appendages such as toes or objects such as styluses may be used to interact with the visualization 138 via the touchscreen. As such, the above gestures should be understood as being illustrative and should not be construed as being limiting in any way. Moreover, the above gestures might find counterparts in other technologies disclosed herein, including those useable by the AR system 144 and/or the VR system 146.

FIG. 1 illustrates one computer system 102, one NVP 104, one SDN network 112, one operating system 114, one NVP application 116, one visualization component 118, one user 136, one visualization 138, one UI 140, one display 142, one AR system 144, one VR system 146, and one input device 148. It should be understood, however, that some implementations of the operating environment 100 can include multiple computer systems 102, multiple NVPs 104, multiple SDN networks 112, multiple operating systems 114, multiple NVP applications 116, multiple visualization components 118, multiple users 136, multiple visualizations 138, multiple UIs 140, multiple displays 142, multiple AR systems 144, multiple VR systems 146, and/or multiple input devices 148. As such, the illustrated embodiment of the operating environment should be understood as being illustrative, and should not be construed as being limiting in any way.

Turning now to FIG. 2, additional aspects of the visualization component 118 will be presented, according to one illustrative embodiment. In particular, FIG. 2 provides further details regarding architecture and subcomponents of the visualization component 118, according to some embodiments. The visualization component 118 can include a number of components and/or subsystems including, but not limited to, an NVP visualization control 200, an NVP visualization engine 202, and/or other components and/or subsystems.

The NVP visualization control 200 can include functionality to receive NVP data input, such as the NVP data 106, from the NVP 104. The NVP visualization control 200 also can include functionality of a visualization user interface used to generate and present the UI 140 described herein above, via the display 142, the AR system 144, and/or the VR system 146.

The NVP visualization engine 202 can include an engine core for generating visualization according to the concepts and technologies described herein, such as the visualization 138. The engine core can utilize any software framework designed for the creation and development of graphics. Some example software frameworks include, but are not limited to, AUTOCAD (available from AUTODESK), UNREAL ENGINE (available from EPIC GAMES, INC.), UNITY (available from UNITY TECHNOLOGIES), CRYENGINE (available from CRYTEK), HAVOK VISION ENGINE (available from HAVOK), other proprietary software frameworks, open source software frameworks, combinations thereof, and the like.

The NVP visualization engine 202 also can include visualization assets for representing application, network link, and/or other resources within visualizations; input and touch modules for interpreting touch and/or multi-touch commands as input; gesture modules for interpreting gesture commands as input; shaders for providing shading of gen-

erated and/or rendered three-dimensional objects; and views functionality for representing and/or interacting with different views of a visualization.

The visualization component **118** also can include various other components and/or subsystems. The visualization component **118** also can include various graphics plugins and/or APIs such as the illustrated DIRECTX APIs, API call emulators, combinations thereof, or the like. It should be appreciated that the visualization component **118** can include additional and/or alternative functionality not shown in FIG. 2. As such, the embodiment illustrated in FIG. 2 should be understood as being illustrative and should not be construed as being limiting in any way.

Turning now to FIG. 3A-3E, diagrams illustrating aspects of various aspects of an example visualization, such as the visualization **138**, will be described, according to embodiments. Turning first to FIG. 3A, a connectivity view **300** shows a plurality of the applications **110** interconnected via at least a portion of the SDN network **112**. The applications **110** are represented as circles of varying gray-scale color. The SDN network **112** is represented as connection lines connecting the applications **110**.

Turning now to FIG. 3B, a full view **302** or service graph view shows a collection of available resources (such as application resources embodied as the applications **110**) along with sequencing and relationships available to be used to process data plane service requests to instantiate a service instance of one or more of the services **108**. A service graph describes the actual topology of the resources instantiated and operationalized by the NVP **104** within the SDN network **112**. The full view **302** also shows a service path **304**. The service path **304** shows a collection of specific resources from the service graph along with sequencing and relationships used to process a specific data plane service request. A service path describes a logical topology of the resources included in a particular instance of a service, such as one of the services **108**. The illustrated service path **304** provided herein follows a flow convention of left-to-right.

A logical view of the service path **304** can be provided by a face plane view **306** of the visualization **138**, such as in the example shown in FIG. 3C. The visualization **138** can provide a plurality of layers of the face plane view **306**. In FIG. 3D, a slice view **308** of the visualization **138** represents one of these layers. Each of the layers can represent a physical technical plant (“TeP”) location within the NVP **104**. A specific resource **310** in a specific TeP location that is processing a particular data plane service request along the service path **304** can be identified by sliding out the corresponding layer for examination by the user **136**, or by navigating into the visualization **138**. A top plane view **312** is shown in FIG. 3E. The top plane view **312** shows the service path **304** across multiple layers of the face plane view **306**—that is, multiple TeP locations within the NVP **104**.

Turning now to FIG. 4, aspects of a method **400** for generating a visualization, such as the visualization **138**, will be described, according to an illustrative embodiment. It should be understood that the operations of the methods disclosed herein are not necessarily presented in any particular order and that performance of some or all of the operations in an alternative order(s) is possible and is contemplated. The operations have been presented in the demonstrated order for ease of description and illustration. Operations may be added, omitted, and/or performed simultaneously, without departing from the scope of the concepts and technologies disclosed herein.

It also should be understood that the methods disclosed herein can be ended at any time and need not be performed in its entirety. Some or all operations of the methods, and/or substantially equivalent operations, can be performed by execution of computer-readable instructions included on a computer storage media, as defined herein. The term “computer-readable instructions,” and variants thereof, as used herein, is used expansively to include routines, applications, application modules, program modules, programs, components, data structures, algorithms, and the like. Computer-readable instructions can be implemented on various system configurations including single-processor or multiprocessor systems, minicomputers, mainframe computers, personal computers, hand-held computing devices, microprocessor-based, programmable consumer electronics, combinations thereof, and the like.

Thus, it should be appreciated that the logical operations described herein are implemented (1) as a sequence of computer implemented acts or program modules running on a computing system and/or (2) as interconnected machine logic circuits or circuit modules within the computing system. The implementation is a matter of choice dependent on the performance and other requirements of the computing system. Accordingly, the logical operations described herein are referred to variously as states, operations, structural devices, acts, or modules. These states, operations, structural devices, acts, and modules may be implemented in software, in firmware, in special purpose digital logic, and any combination thereof. As used herein, the phrase “cause a processor to perform operations” and variants thereof is used to refer to causing a processor of the computer system **102**, a processor of the NVP **104**, such as one or more of the compute resources **126**, and/or a processor one or more other computing systems and/or devices disclosed herein to perform operations.

For purposes of illustrating and describing some of the concepts of the present disclosure, the methods disclosed herein are described as being performed, at least in part, by the computer system **102** via execution of one or more software modules, which may be embodied, at least in part, by the NVP application **116** and the visualization component **118**, described above with reference to FIGS. 1 and 2. It should be understood that additional and/or alternative devices and/or network nodes can provide the functionality described herein via execution of one or more modules, applications, and/or other software. Thus, the illustrated embodiments are illustrative, and should not be viewed as being limiting in any way.

The method **400** will be described with reference to FIG. 4 and further reference to FIGS. 1 and 3. The method **400** begins at operation **402**, where the computer system **102** obtains data associated with the NVP **104**. The data can include the NVP data **106** described above with reference to FIG. 1. Using the NVP data **106**, the computer system **102**, at operation **404**, can generate one or more visualizations of the NVP data **106**, such as the visualization **138**. From operation **404**, the method **400** proceeds to operation **406**, where the computer system **102** outputs the visualization **138** to the display **142**, the AR system **144**, and/or the VR system **146**.

From operation **406**, the method **400** proceeds to operation **408**, where the computer system **102** receives user input from the user **136** and the computer system **102** to interact with the visualization **138** through the UI **140** using various actions. The actions can include user input received by the computer system **102** via the display **142**, the AR system **144**, the VR system **146**, the input device(s) **148**, or some

combination thereof from the user 136. The actions can include movement actions such as, but not limited to, pan up, pan down, pan left, pan right, rotate clockwise, rotate counter-clockwise, tilt down, tilt up, zoom in, zoom out, field-of-view increase, field-of-view decrease, and the like. The movement actions can be used to navigate the visualization 138. Other actions to add, delete, and/or modify at least a portion of the visualization 138 also are contemplated. From operation 408, the method 400 proceeds to operation 410, where the method 400 ends.

Turning now to FIG. 5, a method 500 for adjusting a view of visualizations, such as the visualization 138 will be described in detail, according to an illustrative embodiment. The method 500 will be described from the perspective of the computer system 102. Moreover, the method 500 will be described with reference to FIG. 5 and further reference to FIGS. 1 and 3.

The method 500 begins and proceeds to operation 502, where the computer system 102 receives a user input to adjust a view of the visualization 138. From operation 502, the method 500 proceeds to operation 504, where the computer system 102 generates a new view of the visualization 138 in accordance with the user input. Examples of some illustrative views are provided herein in FIGS. 3A-3E. From operation 504, the method 500 proceeds to operation 506, where the computer system 102 outputs the visualization 138 in the new view. From operation 506, the method 500 proceeds to operation 508, where the method 500 ends.

Turning now to FIG. 6, a method 600 for generating a logical topology view of a visualization, such as the visualization 138, when a new service or a new service instance is instantiated by the NVP 104, will be described in detail, according to an illustrative embodiment. The method 600 will be described from the perspective of the computer system 102. Moreover, the method 600 will be described with reference to FIG. 6 and further reference to FIGS. 1 and 3.

The method 600 begins and proceeds to operation 602, where the computer system 102 receives a request to instantiate a new service instance for one of the services 108. From operation 602, the method 600 proceeds to operation 604, where the computer system 102 generates a logical topology view, such as the face plane view 306 shown in FIG. 3C, for a service path representing the newly-instantiated service instance. From operation 604, the method 600 proceeds to operation 606, where the computer system 102 outputs the logical topology view of the visualization 138. From operation 606, the method 600 proceeds to operation 608, where the method 600 ends.

Turning now to FIG. 7, a method 700 for generating actual topology view of a visualization, such as the visualization 138, in response to a request to utilize a service, such as one of the services 108, will be described in detail, according to an illustrative embodiment. The method 700 will be described from the perspective of the computer system 102. Moreover, the method 700 will be described with reference to FIG. 7 and further reference to FIGS. 1 and 3.

The method 700 begins and proceeds to operation 702, where the computer system 102 receives a request to utilize one of the services 108. From operation 702, the method 700 proceeds to operation 704, where the computer system 102 generates an actual topology view, such as the slice view 308 shown in FIG. 3D. A service graph describes the actual topology of the resources instantiated and operationalized by the NVP 104. From operation 704, the method 700 proceeds to operation 706, where the computer system 102 outputs the

actual topology view of the visualization 138. From operation 706, the method 700 proceeds to operation 708, where the method 700 ends.

Turning now to FIG. 8, a block diagram illustrating a plurality of modules 800 associated with the NVP application 116 and the visualization component 118 that can be executed by the computer system 102 (all best shown in FIG. 1) to perform various operations to facilitate user interaction with a visualization, such as the visualization 138 (shown in FIG. 1), some illustrative examples of which are shown in FIGS. 3A-3E, will be described. The plurality of modules 800 can be executed by the computer system 102 to perform service management operations, and specifically end-to-end service quality management ("E2E SQM") visualization and control. The plurality of modules 800 includes a service definition module 802, a service elements topology assembly module 804, a customer agreement module 806, a resource assignment module 808, a service monitoring module 810, and a visualization and control module 812. The plurality of modules 800 can be interacted with by one or more users, such as the user 136, who might be a service designer, a service engineer, a service salesperson, or the like.

The service definition module 802 can allow a service designer to define, via a user interface such as the user interface 140 (shown in FIG. 1), a service, such as one of the services 108 (also shown in FIG. 1). The service definition module 802 can provide input received from the service designer to the service elements topology assembly module 804. The service definition module 802 can allow a service engineer to identify service elements needed and a topology that chains the service elements in order to effect instances of the service when requested.

Based on overall end-to-end throughput that can be achieved from resources needed to instantiate the service, a dimensioning profile can be established by the resource assignment module 808. The dimensioning profile can set one or more objectives with regard to predictions of end-to-end transaction throughput achieved by certain sizes of resources for typical demand curves to within statistical availability requirements (e.g. five 9's).

Salespeople can utilize the customer agreement module 806 to gain customer's agreement to use the service. The resource assignment module 808 can provision resources and can provide the service based on the customer's anticipated level of demand in accordance with the customer's agreement.

Once instantiated, the service monitoring module 810 can monitor the service for compliance with general objective as well as specific agreements made with particular customers. The visualization and control module 812 can provide the visualization 138. The visualization 138 can provide a way to visualize the currently active topology of an instantiated service, the resource state of the service, and the end-to-end service quality being delivered to the customer. The visualization and control module 812 also can provide the user interface 140 for the control and adjustment of resources assigned by the resource assignment module 808 to provide or augment the service execution, providing a closed-loop control system for adjusting QoS.

FIG. 9 is a block diagram illustrating a computer system 900 configured to provide the functionality in accordance with various embodiments of the concepts and technologies disclosed herein. In some implementations, at least a portion of the hardware resources in the hardware resource layer 120 (best illustrated in FIG. 1) are provided, at least in part, by one or more host server computers (collectively, "host server

cluster”), which is/are configured like the architecture of the computer system **900**. It should be understood, however, that modification to the architecture may be made to facilitate certain interactions among elements described herein. In some implementations, the compute resources **126**, the memory resources **128** and/or the other resources **130** are configured like the architecture of the computer system **900** or portions thereof. In some implementations, the computer system **102**, the AR system **144**, and/or the VR system **146** are configured like the architecture of the computer system **900** or portions thereof.

The computer system **900** includes a processing unit **902**, a memory **904**, one or more user interface devices **906**, one or more input/output (“I/O”) devices **908**, and one or more network devices **910**, each of which is operatively connected to a system bus **912**. The bus **912** enables bi-directional communication between the processing unit **902**, the memory **904**, the user interface devices **906**, the I/O devices **908**, and the network devices **910**.

The processing unit **902** may be a standard central processor that performs arithmetic and logical operations, a more specific purpose programmable logic controller (“PLC”), a programmable gate array, or other type of processor known to those skilled in the art and suitable for controlling the operation of the server computer. Processing units are generally known, and therefore are not described in further detail herein. The compute resources **126** can include one or more instances of the processing units **902**.

The memory **904** communicates with the processing unit **902** via the system bus **912**. In some embodiments, the memory **904** is operatively connected to a memory controller (not shown) that enables communication with the processing unit **902** via the system bus **912**. The memory resources **128** can include one or more instances of the memory **904**. The illustrated memory **904** includes an operating system **914** and one or more program modules **916**.

The operating system **914** can include the operating system **114** of the computer system **102** best shown in FIG. **1**. The operating system **914** can include, but is not limited to, members of the WINDOWS, WINDOWS CE, and/or WINDOWS MOBILE families of operating systems from MICROSOFT CORPORATION, the LINUX family of operating systems, the SYMBIAN family of operating systems from SYMBIAN LIMITED, the BREW family of operating systems from QUALCOMM CORPORATION, the MAC OS, OS X, and/or iOS families of operating systems from APPLE CORPORATION, the FREEBSD family of operating systems, the SOLARIS family of operating systems from ORACLE CORPORATION, other operating systems, and the like.

The program modules **916** may include various software and/or program modules to perform the various operations described herein. For example, the program modules can include the NVP application **116** and the visualization component **118** of the computer system **102** best shown in FIG. **1**, and more particularly, the plurality of modules **800** described above with reference to FIG. **8**. The program modules **916** and/or other programs can be embodied in computer-readable media containing instructions that, when executed by the processing unit **902**, perform various operations such as those described herein. According to embodiments, the program modules **916** may be embodied in hardware, software, firmware, or any combination thereof.

By way of example, and not limitation, computer-readable media may include any available computer storage media or communication media that can be accessed by the computer system **900**. Communication media includes com-

puter-readable instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics changed or set in a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of the any of the above should also be included within the scope of computer-readable media.

Computer storage media includes volatile and non-volatile, removable and non-removable media implemented in any method or technology for storage of information such as computer-readable instructions, data structures, program modules, or other data. Computer storage media includes, but is not limited to, RAM, ROM, EPROM, EEPROM, flash memory or other solid state memory technology, CD-ROM, DVD, or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by the computer system **900**. In the claims, the phrase “computer storage medium” and variations thereof does not include waves or signals per se and/or communication media.

The user interface devices **906** may include one or more devices with which a user accesses the computer system **900**. The user interface devices **906** may include, but are not limited to, computers, servers, personal digital assistants, cellular phones, or any suitable computing devices. The I/O devices **908** enable a user to interface with the program modules **916**. In one embodiment, the I/O devices **908** are operatively connected to an I/O controller (not shown) that enables communication with the processing unit **902** via the system bus **912**. The I/O devices **908** may include one or more input devices, such as, but not limited to, a keyboard, a mouse, or an electronic stylus. Further, the I/O devices **908** may include one or more output devices, such as, but not limited to, a display screen or a printer.

The network devices **910** enable the computer system **900** to communicate with other networks or remote systems via a network **918**, which can include the SDN network **112** and/or other networks described herein. Examples of the network devices **910** include, but are not limited to, a modem, a radio frequency (“RF”) or infrared (“IR”) transceiver, a telephonic interface, a bridge, a router, or a network card. The network **918** may include a wireless network such as, but not limited to, a wireless local area network (“WLAN”), a wireless wide area network (“WWAN”), a wireless personal area network (“WPAN”) such as provided via BLUETOOTH technology, a wireless metropolitan area network (“WMAN”) such as a WiMAX network or metropolitan cellular network. Alternatively, the network **918** may be a wired network such as, but not limited to, a wide area network (“WAN”), a wired LAN such as provided via Ethernet, a wired personal area network (“PAN”), or a wired metropolitan area network (“MAN”).

Turning now to FIG. **10**, additional details of a network **1000**, such as the network **108**, are illustrated, according to an illustrative embodiment. The network **1000** includes a cellular network **1002**, a packet data network **1004**, for example, the Internet, and a circuit switched network **1006**, for example, a publicly switched telephone network (“PSTN”). The cellular network **1002** includes various components such as, but not limited to, base transceiver stations (“BTSs”), Node-B’s or e-Node-B’s, base station controllers

(“BSCs”), radio network controllers (“RNCs”), mobile switching centers (“MSCs”), mobile management entities (“MMEs”), short message service centers (“SMSCs”), multimedia messaging service centers (“MMSCs”), home location registers (“HLRs”), home subscriber servers (“HSSs”), visitor location registers (“VLRs”), charging platforms, billing platforms, voicemail platforms, GPRS core network components, location service nodes, an IP Multimedia Subsystem (“IMS”), and the like. The cellular network **1002** also includes radios and nodes for receiving and transmitting voice, data, and combinations thereof to and from radio transceivers, networks, the packet data network **1004**, and the circuit switched network **1006**.

A mobile communications device **1008**, such as, for example, the computer system **102**, a cellular telephone, a user equipment, a mobile terminal, a PDA, a laptop computer, a handheld computer, and combinations thereof, can be operatively connected to the cellular network **1002**. The cellular network **1002** can be configured as a 2G Global System for Mobile communications (“GSM”) network and can provide data communications via General Packet Radio Service (“GPRS”) and/or Enhanced Data rates for GSM Evolution (“EDGE”). Additionally, or alternatively, the cellular network **1002** can be configured as a 3G Universal Mobile Telecommunications System (“UMTS”) network and can provide data communications via the High-Speed Packet Access (“HSPA”) protocol family, for example, High-Speed Downlink Packet Access (“HSDPA”), Enhanced UpLink (“EUL”) (also referred to as High-Speed Uplink Packet Access (“HSUPA”)), and HSPA+. The cellular network **1002** also is compatible with 4G mobile communications standards such as Long-Term Evolution (“LTE”), or the like, as well as evolved and future mobile standards.

The packet data network **1004** includes various devices, for example, servers, computers, databases, and other devices in communication with one another, as is generally known. The packet data network **1004** devices are accessible via one or more network links. The servers often store various files that are provided to a requesting device such as, for example, a computer, a terminal, a smartphone, or the like. Typically, the requesting device includes software (a “browser”) for executing a web page in a format readable by the browser or other software. Other files and/or data may be accessible via “links” in the retrieved files, as is generally known. In some embodiments, the packet data network **1004** includes or is in communication with the Internet. The circuit switched network **1006** includes various hardware and software for providing circuit switched communications. The circuit switched network **1006** may include, or may be, what is often referred to as a plain old telephone system (“POTS”). The functionality of a circuit switched network **1006** or other circuit-switched network are generally known and will not be described herein in detail.

The illustrated cellular network **1002** is shown in communication with the packet data network **1004** and a circuit switched network **1006**, though it should be appreciated that this is not necessarily the case. One or more Internet-capable devices **1010**, for example, the computer system **102**, a PC, a laptop, a portable device, or another suitable device, can communicate with one or more cellular networks **1002**, and devices connected thereto, through the packet data network **1004**. It also should be appreciated that the Internet-capable device **1010** can communicate with the packet data network **904** through the circuit switched network **1006**, the cellular network **1002**, and/or via other networks (not illustrated).

As illustrated, a communications device **1012**, for example, a telephone, facsimile machine, modem, computer, or the like, can be in communication with the circuit switched network **1006**, and therethrough to the packet data network **1004** and/or the cellular network **1002**. It should be appreciated that the communications device **1012** can be an Internet-capable device, and can be substantially similar to the Internet-capable device **1010**. In the specification, the network **1000** is used to refer broadly to any combination of the networks **1002**, **1004**, **1006**. It should be appreciated that substantially all of the functionality described with reference to the network **1000** can be performed by the cellular network **1002**, the packet data network **1004**, and/or the circuit switched network **1006**, alone or in combination with other networks, network elements, and the like.

Turning now to FIG. **11**, an illustrative mobile device **1100** and components thereof will be described. In some embodiments, the computer system **102**, the AR system **144**, and/or the VR system **146** described above with reference to FIG. **1** can be configured as and/or can have an architecture similar or identical to the mobile device **1100** described herein in FIG. **11**. It should be understood, however, that the computer system **102**, the AR system **144**, and/or the VR system **146** may or may not include the functionality described herein with reference to FIG. **11**. While connections are not shown between the various components illustrated in FIG. **11**, it should be understood that some, none, or all of the components illustrated in FIG. **11** can be configured to interact with one other to carry out various device functions. In some embodiments, the components are arranged so as to communicate via one or more busses (not shown). Thus, it should be understood that FIG. **11** and the following description are intended to provide a general understanding of a suitable environment in which various aspects of embodiments can be implemented, and should not be construed as being limiting in any way.

As illustrated in FIG. **11**, the mobile device **1100** can include a display **1102** for displaying data. According to various embodiments, the display **1102** can be configured to display the visualization **138**, various graphical user interface (“GUI”) elements, text, images, video, advertisements, prompts, virtual keypads and/or keyboards, messaging data, notification messages, metadata, internet content, device status, time, date, calendar data, device preferences, map and location data, combinations thereof, and the like. The mobile device **1100** also can include a processor **1104** and a memory or other data storage device (“memory”) **1106**. The processor **1104** can be configured to process data and/or can execute computer-executable instructions stored in the memory **1106**. The computer-executable instructions executed by the processor **1104** can include, for example, an operating system **1108** (e.g., the operating system **114**), one or more applications **1110** (e.g., the NVP application **116** and/or the visualization component **118**), other computer-executable instructions stored in a memory **1106**, or the like. In some embodiments, the applications **1110** also can include a UI application (not illustrated in FIG. **11**).

The UI application (e.g., the UI **140**) can interface with the operating system **1108** to facilitate user interaction with functionality and/or data stored at the mobile device **1100** and/or stored elsewhere. In some embodiments, the operating system **1108** can include a member of the SYMBIAN OS family of operating systems from SYMBIAN LIMITED, a member of the WINDOWS MOBILE OS and/or WINDOWS PHONE OS families of operating systems from MICROSOFT CORPORATION, a member of the PALM WEBOS family of operating systems from HEWLETT

PACKARD CORPORATION, a member of the BLACKBERRY OS family of operating systems from RESEARCH IN MOTION LIMITED, a member of the IOS family of operating systems from APPLE INC., a member of the ANDROID OS family of operating systems from GOOGLE INC., and/or other operating systems. These operating systems are merely illustrative of some contemplated operating systems that may be used in accordance with various embodiments of the concepts and technologies described herein and therefore should not be construed as being limiting in any way.

The UI application can be executed by the processor **1104** to aid a user, such as the user **136**, in viewing and interacting with the visualization **138**, entering content, viewing account information, answering/initiating calls, entering/deleting data, entering and setting user IDs and passwords for device access, configuring settings, manipulating address book content and/or settings, multimode interaction, interacting with other applications **1110**, and otherwise facilitating user interaction with the operating system **1108**, the applications **1110**, and/or other types or instances of data **1112** that can be stored at the mobile device **1100**.

According to various embodiments, the applications **1110** can include, for example, presence applications, visual voice mail applications, messaging applications, text-to-speech and speech-to-text applications, add-ons, plug-ins, email applications, music applications, video applications, camera applications, location-based service applications, power conservation applications, game applications, productivity applications, entertainment applications, enterprise applications, combinations thereof, and the like. The applications **1110**, the data **1112**, and/or portions thereof can be stored in the memory **1106** and/or in a firmware **1114**, and can be executed by the processor **1104**. The firmware **1114** also can store code for execution during device power up and power down operations. It can be appreciated that the firmware **1114** can be stored in a volatile or non-volatile data storage device including, but not limited to, the memory **1106** and/or a portion thereof.

The mobile device **1100** also can include an input/output (“I/O”) interface **1116**. The I/O interface **1116** can be configured to support the input/output of data such as location information, user information, organization information, presence status information, user IDs, passwords, and application initiation (start-up) requests. In some embodiments, the I/O interface **1116** can include a hardwire connection such as USB port, a mini-USB port, a micro-USB port, an audio jack, a PS2 port, an IEEE 1394 (“FIREWIRE”) port, a serial port, a parallel port, an Ethernet (RJ45) port, an RJ11 port, a proprietary port, combinations thereof, or the like. In some embodiments, the mobile device **1100** can be configured to synchronize with another device to transfer content to and/or from the mobile device **1100**. In some embodiments, the mobile device **1100** can be configured to receive updates to one or more of the applications **1110** via the I/O interface **1116**, though this is not necessarily the case. In some embodiments, the I/O interface **1116** accepts I/O devices such as the input device(s) **148**, keyboards, keypads, mice, interface tethers, printers, plotters, external storage, touch/multi-touch screens, touch pads, trackballs, joysticks, microphones, remote control devices, displays, projectors, medical equipment (e.g., stethoscopes, heart monitors, and other health metric monitors), modems, routers, external power sources, docking stations, combinations thereof, and the like. It should be appreciated that the I/O interface **1116** may be used for communications between the mobile device **1100** and a network device or local device.

The mobile device **1100** also can include a communications component **1118**. The communications component **1118** can be configured to interface with the processor **1104** to facilitate wired and/or wireless communications with one or more networks described above herein. In some embodiments, other networks include networks that utilize non-cellular wireless technologies such as WI-FI or WIMAX. In some embodiments, the communications component **1118** includes a multimode communications subsystem for facilitating communications via the cellular network and one or more other networks.

The communications component **1118**, in some embodiments, includes one or more transceivers. The one or more transceivers, if included, can be configured to communicate over the same and/or different wireless technology standards with respect to one another. For example, in some embodiments one or more of the transceivers of the communications component **1118** may be configured to communicate using GSM, cod division multiple access (“CDMA”), CDMAONE, CDMA2000, LTE, and various other 2G, 2.5G, 3G, 4G, and greater generation technology standards. Moreover, the communications component **1118** may facilitate communications over various channel access methods (which may or may not be used by the aforementioned standards) including, but not limited to, time division multiple access (“TDMA”), frequency division multiple access (“FDMA”), wideband CDMA (“W-CDMA”), orthogonal frequency-division multiplexing (“OFDM”), spatial division multiple access (“SDMA”), and the like.

In addition, the communications component **1118** may facilitate data communications using GPRS, EDGE, the HSPA protocol family, including HSDPA, EUL, or otherwise termed HSUPA, HSPA+, and various other current and future wireless data access standards. In the illustrated embodiment, the communications component **1118** can include a first transceiver (“TxRx”) **1120A** that can operate in a first communications mode (e.g., GSM). The communications component **1118** also can include an N<sup>th</sup> transceiver (“TxRx”) **1120N** that can operate in a second communications mode relative to the first transceiver **1120A** (e.g., UMTS). While two transceivers **1120A-N** (hereinafter collectively and/or generically referred to as “transceivers **1120**”) are shown in FIG. **11**, it should be appreciated that less than two, two, and/or more than two transceivers **1120** can be included in the communications component **1118**.

The communications component **1118** also can include an alternative transceiver (“Alt TxRx”) **1122** for supporting other types and/or standards of communications. According to various contemplated embodiments, the alternative transceiver **1122** can communicate using various communications technologies such as, for example, WI-FI, WIMAX, BLUETOOTH, infrared, infrared data association (“IRDA”), near-field communications (“NFC”), other radio frequency (“RF”) technologies, combinations thereof, and the like.

In some embodiments, the communications component **1118** also can facilitate reception from terrestrial radio networks, digital satellite radio networks, internet-based radio service networks, combinations thereof, and the like. The communications component **1118** can process data from a network such as the Internet, an intranet, a broadband network, a WI-FI hotspot, an Internet service provider (“ISP”), a digital subscriber line (“DSL”) provider, a broadband provider, combinations thereof, or the like.

The mobile device **1100** also can include one or more sensors **1124**. The sensors **1124** can include temperature sensors, light sensors, air quality sensors, movement sen-

sors, orientation sensors, noise sensors, proximity sensors, or the like. As such, it should be understood that the sensors **1124** can include, but are not limited to, accelerometers, magnetometers, gyroscopes, infrared sensors, noise sensors, microphones, combinations thereof, or the like. Additionally, audio capabilities for the mobile device **1100** may be provided by an audio I/O component **1126**. The audio I/O component **1126** of the mobile device **1100** can include one or more speakers for the output of audio signals, one or more microphones for the collection and/or input of audio signals, and/or other audio input and/or output devices.

The illustrated mobile device **1100** also can include a subscriber identity module (“SIM”) system **1128**. The SIM system **1128** can include a universal SIM (“USIM”), a universal integrated circuit card (“UICC”) and/or other identity devices. The SIM system **1128** can include and/or can be connected to or inserted into an interface such as a slot interface **1130**. In some embodiments, the slot interface **1130** can be configured to accept insertion of other identity cards or modules for accessing various types of networks. Additionally, or alternatively, the slot interface **1130** can be configured to accept multiple subscriber identity cards. Because other devices and/or modules for identifying users and/or the mobile device **1100** are contemplated, it should be understood that these embodiments are illustrative, and should not be construed as being limiting in any way.

The mobile device **1100** also can include an image capture and processing system **1132** (“image system”). The image system **1132** can be configured to capture or otherwise obtain photos, videos, and/or other visual information. As such, the image system **1132** can include cameras, lenses, charge-coupled devices (“CCDs”), combinations thereof, or the like. The mobile device **1100** may also include a video system **1134**. The video system **1134** can be configured to capture, process, record, modify, and/or store video content. Photos and videos obtained using the image system **1132** and the video system **1134**, respectively, may be added as message content to a multimedia message service (“MMS”) message, email message, and sent to another mobile device. The video and/or photo content also can be shared with other devices via various types of data transfers via wired and/or wireless communication devices as described herein.

The mobile device **1100** also can include one or more location components **1136** (e.g., the location component(s) **118**). The location components **1136** can be configured to send and/or receive signals to determine a geographic location of the mobile device **1100**. According to various embodiments, the location components **1136** can send and/or receive signals from GPS devices, A-GPS devices, WI-FI/WIMAX and/or cellular network triangulation data, combinations thereof, and the like. The location component **1136** also can be configured to communicate with the communications component **1118** to retrieve triangulation data for determining a location of the mobile device **1100**. In some embodiments, the location component **1136** can interface with cellular network nodes, telephone lines, satellites, location transmitters and/or beacons, wireless network transmitters and receivers, combinations thereof, and the like. In some embodiments, the location component **1136** can include and/or can communicate with one or more of the sensors **1124** such as a compass, an accelerometer, and/or a gyroscope to determine the orientation of the mobile device **1100**. Using the location component **1136**, the mobile device **1100** can generate and/or receive data to identify its geographic location, or to transmit data used by other devices to determine the location of the mobile device **1100**. The

location component **1136** may include multiple components for determining the location and/or orientation of the mobile device **1100**.

The illustrated mobile device **1100** also can include a power source **1138**. The power source **1138** can include one or more batteries, power supplies, power cells, and/or other power subsystems including alternating current (“AC”) and/or direct current (“DC”) power devices. The power source **1138** also can interface with an external power system or charging equipment via a power I/O component **1140**. Because the mobile device **1100** can include additional and/or alternative components, the above embodiment should be understood as being illustrative of one possible operating environment for various embodiments of the concepts and technologies described herein. The described embodiment of the mobile device **1100** is illustrative, and should not be construed as being limiting in any way.

Based on the foregoing, it should be appreciated that concepts and technologies directed to visualizations for NVPs have been disclosed herein. Although the subject matter presented herein has been described in language specific to computer structural features, methodological and transformative acts, specific computing machinery, and computer-readable media, it is to be understood that the concepts and technologies disclosed herein are not necessarily limited to the specific features, acts, or media described herein. Rather, the specific features, acts and mediums are disclosed as example forms of implementing the concepts and technologies disclosed herein.

The subject matter described above is provided by way of illustration only and should not be construed as limiting. Various modifications and changes may be made to the subject matter described herein without following the example embodiments and applications illustrated and described, and without departing from the true spirit and scope of the embodiments of the concepts and technologies disclosed herein.

We claim:

1. A method comprising:

obtaining, at a computer system comprising a processor that is executing a visualization component, network virtualization platform data associated with a network virtualization platform that decouples hardware and software for a software defined network;

generating, by the computer system executing the visualization component, a three-dimensional visualization of the network virtualization platform data associated with the network virtualization platform,

wherein the three-dimensional visualization comprises a three-dimensional central cube visualization that comprises a plurality of layers with each layer corresponding with a physical technical plant location within the network virtualization platform,

wherein each of the plurality of layers comprises a plurality of smaller cube virtualizations that each comprise a node representative of at least one application of a plurality of applications executable by at least a portion of a plurality of resources of the network virtualization platform at the physical technical plant location corresponding to one of the plurality of layers,

wherein the three-dimensional central cube visualization provides a service path across the plurality of smaller cube virtualizations so as to indicate a sequence of applications available to process a data

25

plane service request to instantiate a service instance for a service within the software defined network; and  
 outputting, by the computer system, the three-dimensional visualization to present the three-dimensional visualization to a user for display and interaction.

2. The method of claim 1, wherein the three-dimensional visualization is output to at least one of an augmented reality system or a virtual reality system that communicates with the computer system.

3. The method of claim 1, wherein each instance of a layer of the plurality of layers can identify a resource along the service path that is processing the data plane service request at the physical technical plant location by sliding out the layer from the three-dimensional central cube visualization.

4. The method of claim 1, wherein the three-dimensional central cube visualization further comprises a plurality of networking links between at least a portion of a plurality of applications of the plurality of smaller cube virtualizations.

5. The method of claim 4, wherein the network virtualization platform data comprises data associated with a topology of the plurality of resources of the network virtualization platform.

6. The method of claim 5, further comprising:  
 receiving, by the computer system, the data plane service request to instantiate the service instance of the service of the software defined network;  
 in response to receiving the data plane service request, generating, by the computer system, a logical topology view of the three-dimensional central cube visualization; and  
 outputting, by the computer system, the logical topology view of the three-dimensional central cube visualization;  
 wherein the logical topology view comprises at least a portion of the topology of the plurality of resources of the network virtualization platform; and  
 wherein the logical topology view provides the service path comprising a set of the plurality of resources utilized by the service.

7. The method of claim 6, wherein the logical topology view comprises a face view of the three-dimensional central cube visualization, the face view comprising the plurality of layers, and wherein each layer of the plurality of layers comprised in the face view represents one physical technical plant location and a subset of the set of the plurality of resources utilized by the service within the physical technical plant location.

8. The method of claim 5, further comprising:  
 receiving, by the computer system, a request to utilize the service;  
 generating, by the computer system, an actual topology view of at least a portion of the topology of the plurality of resources of the network virtualization platform, wherein the actual topology view represents a service graph comprising a set of the plurality of resources utilized by the service; and  
 outputting, by the computer system, the actual topology view.

9. A computer system comprising:  
 a processor; and  
 a memory comprising instructions for a visualization component that, when executed by the processor, cause the processor to perform operations comprising:

26

obtaining network virtualization platform data associated with a network virtualization platform that decouples hardware and software for a software defined network,  
 generating a three-dimensional visualization of the network virtualization platform data associated with the network virtualization platform,  
 wherein the three-dimensional visualization comprises a three-dimensional central cube visualization that comprises a plurality of layers with each layer corresponding with a physical technical plant location within the network virtualization platform,  
 wherein each of the plurality of layers comprises a plurality of smaller cube virtualizations that each comprise a node representative of at least one application of a plurality of applications executable by at least a portion of a plurality of resources of the network virtualization platform at the physical technical plant location corresponding to one of the plurality of layers,  
 wherein the three-dimensional central cube visualization provides a service path across the plurality of smaller cube virtualizations so as to indicate a sequence of applications available to process a data plane service request to instantiate a service instance for a service within the software defined network, and  
 outputting the three-dimensional visualization to present the three-dimensional visualization to a user for display and interaction.

10. The computer system of claim 9, wherein the three-dimensional visualization is output to at least one of an augmented reality system or a virtual reality system that communicates with the computer system.

11. The computer system of claim 9, wherein each instance of a layer of the plurality of layers can identify a resource along the service path that is processing the data plane service request at the physical technical plant location by sliding out the layer from the three-dimensional central cube visualization.

12. The computer system of claim 11, wherein the three-dimensional central cube visualization further comprises a plurality of networking links between at least a portion of the plurality of applications of the plurality of smaller cube virtualizations.

13. The computer system of claim 12, wherein the network virtualization platform data comprises data associated with a topology of the plurality of resources of the network virtualization platform.

14. The computer system of claim 13, wherein the operations further comprise:  
 receiving the data plane service request to instantiate the service instance of the service of the software defined network;  
 in response to receiving the data plane service request, generating a logical topology view of the three-dimensional central cube visualization; and  
 outputting the logical topology view of the three-dimensional central cube visualization;  
 wherein the logical topology view comprises at least a portion of the topology of the plurality of resources of the network virtualization platform; and  
 wherein the logical topology view provides the service path comprising a set of the plurality of resources utilized by the service.

27

15. The computer system of claim 14, wherein the logical topology view comprises a face view of the three-dimensional central cube visualization, the face view comprising the plurality of layers, and wherein each layer of the plurality of layers comprised in the face view represents one physical technical plant location and a subset of the set of the plurality of resources utilized by the service within the physical technical plant location.

16. The computer system of claim 13, wherein the operations further comprise:

- receiving a request to utilize the service;
- generating an actual topology view of at least a portion of the topology of the plurality of resources of the network virtualization platform, wherein the actual topology view represents a service graph comprising a set of the plurality of resources utilized by the service; and
- outputting the actual topology view.

17. A computer-readable storage medium having instructions associated with a visualization component stored thereon that, when executed by a computer system, cause the computer system to perform operations comprising:

- obtaining network virtualization platform data associated with a network virtualization platform that decouples hardware and software for a software defined network;
- generating a three-dimensional visualization of the network virtualization platform data associated with the network virtualization platform, wherein the three-dimensional visualization comprises a three-dimensional central cube visualization that comprises a plurality of layers with each layer corresponding with a physical technical plant location within the network virtualization platform,
- wherein each of the plurality of layers comprises a plurality of smaller cube virtualizations, each of the smaller cube virtualizations comprising a node representative of an application executable by at least a portion of a plurality of resources of the network virtualization platform at the physical technical plant location corresponding to one of the plurality of layers,

28

wherein the three-dimensional central cube visualization provides a service path across the plurality of smaller cube virtualizations so as to indicate a sequence of applications that are available to process a data plane service request to instantiate a service instance for a service within the software defined network; and

outputting the three-dimensional visualization to present the three-dimensional visualization to a user for display and interaction.

18. The computer-readable storage medium of claim 17, wherein the network virtualization platform data is associated with a topology of the plurality of resources of the network virtualization platform.

19. The computer-readable storage medium of claim 18, wherein the operations further comprise:

- receiving the data plane service request to instantiate the service instance of the service;
- in response to receiving the data plane service request, generating a logical topology view of the three-dimensional central cube visualization; and
- outputting the logical topology view of the three-dimensional central cube visualization;

wherein the logical topology view comprises at least a portion of the topology of the plurality of resources of the network virtualization platform; and

wherein the logical topology view represents the service path comprising a set of the plurality of resources utilized by the service.

20. The computer-readable storage medium of claim 18, wherein the operations further comprise:

- receiving a request to utilize the service;
- generating an actual topology view of at least a portion of the topology of the plurality of resources of the network virtualization platform, wherein the actual topology view represents a service graph comprising a set of the plurality of resources utilized by the service; and
- outputting the actual topology view.

\* \* \* \* \*